

HB31
M415
no.09-16



DEWEY 11

Massachusetts Institute of Technology
Department of Economics
Working Paper Series

Noisy Business Cycles

George-Marios Angeletos
Jennifer La'O

Working Paper 09-16
May 14, 2009

Room E52-251
50 Memorial Drive
Cambridge, MA 02142

This paper can be downloaded without charge from the
Social Science Research Network Paper Collection at
<http://ssrn.com/abstract=1414302>

Noisy Business Cycles*

George-Marios Angeletos

MIT and NBER

Jennifer La'O

MIT

May 14, 2009


Abstract

This paper investigates a real-business-cycle economy that features dispersed information about the underlying aggregate productivity shocks, taste shocks, and—potentially—shocks to monopoly power. We show how the dispersion of information can (i) contribute to significant inertia in the response of macroeconomic outcomes to such shocks; (ii) induce a negative short-run response of employment to productivity shocks; (iii) imply that productivity shocks explain only a small fraction of high-frequency fluctuations; (iv) contribute to significant noise in the business cycle; (v) formalize a certain type of demand shocks within an RBC economy; and (vi) generate cyclical variation in observed Solow residuals and labor wedges. Importantly, none of these properties requires significant uncertainty about the underlying fundamentals: they rest on the heterogeneity of information and the strength of trade linkages in the economy, not the level of uncertainty. Finally, none of these properties are symptoms of inefficiency: apart from undoing monopoly distortions or providing the agents with more information, no policy intervention can improve upon the equilibrium allocations.

JEL codes: C7, D6, D8.

Keywords: Business cycles, fluctuations, heterogeneous information, informational frictions, noise, strategic complementarity, higher-order beliefs.

*This paper was prepared for the 2009 NBER Macroeconomics Annual. We are grateful for their detailed feedback to our discussants, Christian Hellwig and Robert King, and the organizers, Daron Acemoglu, Kenneth Rogoff, and Michael Woodford. We also received useful comments from Philippe Bacchetta, Ricardo Caballero, V.V. Chari, Emmanuel Farhi, Mikhail Golosov, Guido Lorenzoni, Ellen McGrattan, Patrick Kehoe, Ricardo Lagos, Stephen Morris, Alessandro Pavan, Ricardo Reis, Robert Shimer, Robert Townsend, and seminar participants at MIT, the University of Wisconsin at Madison, the Federal Reserve Bank of Minneapolis, the Federal Reserve Board, the World Bank, the 2009 conference in honor of Truman Bewley at the University of Texas, Austin, and the 2009 NBER Macroeconomics Annual. *Email:* angelet@mit.edu, jenlao@mit.edu.



Digitized by the Internet Archive
in 2011 with funding from
Boston Library Consortium Member Libraries

<http://www.archive.org/details/noisybusinesscyc00ange>

1 Introduction

There is a long tradition in macroeconomics, going back to Phelps (1970), Lucas (1972, 1975), Barro (1976), King (1982), and others, to use information frictions to break the neutrality of monetary policy. This literature has recently been revived by Mankiw and Reis (2002), Sims (2003), Woodford (2003a, 2008), Mackowiak and Wiederholt (2008), and others. While this work has proposed new formalizations of the origins of informational frictions, most of it has remained focused on the old theme of breaking monetary neutrality. In this paper, we are also concerned with informational frictions, but shift focus to a very different theme: we study how informational frictions impact the response of the economy to aggregate productivity shocks, and other *real* shocks, within the context of a micro-founded real-business-cycle model.

This shift is motivated by two considerations. First, the empirical relevance of theories that require significant lack of information, or some type of unawareness, about the current monetary policy is debatable. Indeed, the first generation of the aforementioned literature succumbed to the criticism that such information is widely, readily, and cheaply available.¹ Second, we contend that the dispersion of information about the *real* shocks hitting the economy is far more severe than the one about the conduct of monetary policy. In the ongoing crisis, for example, there is far more uncertainty, and disagreement, about non-monetary factors such as the value of certain assets, the health of the financial system, or the broader economic fundamentals. And yet, the pertinent literature has little to say about how the heterogeneity of information about the real underlying economic fundamentals matters for macroeconomic outcomes. Finally, we would like to understand not only how this heterogeneity may impact the positive properties of the business cycle but also its normative properties.

Motivated by these considerations, this paper introduces dispersed information in an otherwise canonical RBC model, where nominal prices are flexible and monetary factors are irrelevant. We first show that the dispersion of information can significantly alter certain positive properties of the RBC paradigm—indeed in ways that might imply that technology shocks explain only a small fraction of high-frequency business cycles, while at the same time helping overcome certain criticisms that New-Keynesians have raised against the RBC paradigm. We next show that this significant

¹The new generation has attempted to escape this criticism by postulating that, even if such information is readily and cheaply available, it may still be hard to update one's information sufficiently frequently (Mankiw and Reis, 2002) or to process and absorb such information sufficiently well (Woodford, 2003a, 2008; Mackowiak and Wiederholt, 2008). While these ideas are plausible, it seems hard either to gauge their quantitative importance or to reconcile them with the fact that financial markets respond nearly instantaneously to any news about monetary policy, or that a variety of economic agents appear to pay close attention to monetary policy.

change in the positive properties of the RBC paradigm happens without affecting one important normative lesson: as long as there are no monopoly distortions, the equilibrium allocations coincide with the solution to a certain planning problem, leaving no room for stabilization policies.

These results should not be interpreted narrowly only as an attack against the New-Keynesian paradigm. Our primary goal is to provide a clean theoretical benchmark for the positive and normative implications of dispersed information. Abstracting from nominal frictions best serves this purpose. And yet, our framework is rich enough to nest the real backbone of New-Keynesian models. Our framework and results may thus prove equally useful for RBC and New-Keynesian analysts alike. In this regard, we believe that our paper makes not only a specific contribution into business-cycle theory but also a broader methodological contribution.

Preview of model. The backbone of our model is a canonical RBC economy. We abstract from capital to simplify the analysis, but allow for a continuum of differentiated commodities. This multi-good (or multi-sector) specification serves two purposes. First and foremost, it introduces a certain type of trade interactions that, as first highlighted in Angeletos and La'O (2009b), play a crucial role for aggregate fluctuations when, and only when, information is dispersed; this is true whether each of the goods is produced in a competitive or monopolistic fashion. Second, when combined with monopoly power, this specification permits us to nest the real backbone of New-Keynesian models, facilitating a translation of our results to such models.² Accordingly, while the core of our analysis focuses on shocks to technology (TFP), in principle we also allow for two other types of shocks to the fundamentals of the economy: taste shocks (shocks to the disutility of labor), and mark-up shocks (shocks to the elasticity of demand). However, none of our results rests on the presence of either monopoly power or these additional shocks.

The only friction featured in our model is that certain economic decisions have to be made under heterogeneous information about the aggregate shocks hitting the economy. The challenge is to incorporate this informational friction without an undue sacrifice in either the micro-foundations or the tractability of the analysis. Towards this goal, we formalize this friction with a certain geographical segmentation, following similar lines as Lucas (1972), Barro (1976), Townsend (1983), and Angeletos and La'O (2008, 2009b). In particular, we assume that each period firms and workers meet in different "islands" and have to make their employment and production decisions while facing uncertainty about the shocks hitting other islands. At the same time, we assume that consumption choices take place in a centralized market, where information is homogenous, and that households are "big families", with fully diversified sources of income. This guarantees that our economy admits

²Indeed, all the results we document in this paper directly extend to a New-Keynesian variant as long as monetary policy replicates flexible-price allocations, which in certain cases is the optimal thing to do (Angeletos and Lao, 2008).

a representative consumer and maintains high tractability in analysis despite the fact that some key economic decisions take place under heterogeneous information.

Positive results. As mentioned, the core of our analysis focuses on the special case where firms are competitive and the only shocks hitting the fundamentals of the economy are technology (TFP) shocks which makes the analysis directly comparable to the RBC paradigm.

(i) In standard RBC models (e.g., Hansen, 1985; Prescott, 1986), macroeconomic outcomes respond fast and strongly to technology shocks. We show that the dispersion of information induces inertia in the response of macroeconomic outcomes. Perhaps paradoxically, this inertia can be significant even if the agents face little uncertainty about the underlying shocks.

(ii) Some researchers have argued that employment responds negatively to productivity shocks in the data; have pointed out that that this fact is inconsistent with standard RBC models; and have used this fact to argue in favor of New-Keynesian models (e.g., Galí, 1999; Basu, Fernald and Kimball, 2006; Galí and Rabanal, 2004). Although this fact remains debatable (e.g., Christiano, Eichenbaum and Vigfusson, 2003; McGrattan, 2004), we show that the dispersion of information can accommodate it within the RBC paradigm.

(iii) In the RBC paradigm, technology shocks account for the bulk of short-run fluctuations. Many economists have argued that this is empirically implausible and have favored New-Keynesian alternatives. We show that the dispersion of information can induce technology shocks to explain only a small fraction of the high-frequency variation in the business cycle. And yet, the entire business cycle remains neoclassical in its nature: monetary factors play no role whatsoever.

(iv) What drives the residual variation in short-run fluctuations in our model is simply the noise in available information, that is, correlated errors in the agents' expectations of the underlying technology shocks. Most interestingly, we show that the fraction of short-run volatility that is due to such noise can be arbitrarily high even if the agents are nearly perfectly informed about the underlying technology shocks.

(v) These noise-driven fluctuations help formalize a certain type of "demand shocks" within an RBC setting. The associated errors in forecasting economic activity can be interpreted as variation in expectations of "aggregate demand". They help increase the relative volatility of employment while decreasing its correlation with output. An identification strategy as in Blanchard and Quah (1989) or Galí (1999) would likely identify these shocks as "demand" shocks.

(vi) These noise-driven fluctuations involve countercyclical variation in measured labor wedges, and procyclical variation in Solow residuals, consistent with what observed in the data. Once again, these cyclical variations can be significant even if the agents are nearly perfectly informed about the underlying technology shocks.

While we stop short of quantifying these results, we hope that they at least highlight how the heterogeneity of information has a very different mark on macroeconomics outcomes than the uncertainty about fundamentals—a point that we further elaborate on in Angeletos and La’O (2009b). Indeed, what drives our results is not per se the level of uncertainty about the underlying technology or other shocks, but rather the lack of common knowledge about them: our effects are consistent with an arbitrarily small level of uncertainty about the underlying fundamentals.

At the same time, the lack of common knowledge does not *alone* explain the magnitude of our effects. Rather, this depends crucially on the strength of trade linkages among the firms and workers in our economy. This idea is formalized by our game-theoretic representation. A measure of the trade linkages in our economy, namely the elasticity of substitution across different goods, maps one-to-one to the degree of strategic complementarity in the game that represents our economy. One can then extrapolate from earlier more abstract work on games of strategic complementarity (Morris and Shin, 2002, Angeletos and Pavan, 2007a) that the strength of trade linkages in our economy may play a crucial role in determining the equilibrium effects of heterogeneous information. We conclude that our findings hinge on the combination of heterogeneous information with strong trade linkages—but they do not hinge on the level of uncertainty about the underlying fundamentals.

We finally seek to understand the normative content of the aforementioned findings. We know that a planner could improve welfare by aggregating the information that is dispersed in the economy, or otherwise providing the agents with more information. But this provides no guidance on whether the government should stabilize the fluctuations that originate in noise, or otherwise interfere with the way the economy responds to available information. To address this issue, one has to ask whether a planner can improve upon the equilibrium allocations *without* changing the information structure.

We show that the answer to this question is essentially negative. In particular, in the special case of our model where firms are competitive, there is indeed no way in which the planner can raise welfare without changing the information that is available to the economy. As for the more general case where firms have monopoly power, the best the planner can do is merely to undo the monopoly distortion, much alike what he is supposed to do when information is commonly shared. We conclude that, insofar the information is taken as exogenous, the key normative lessons of the pertinent business-cycle theory survive the introduction of dispersed information, no matter how severely the positive lessons might be affected.

Layout. The remainder of the introduction discusses the related literature. Section 2 introduces the model. Section 3 characterizes the general equilibrium. Sections 4 and 5 explore the implications for business cycles. Section 6 studies efficiency. Section 7 concludes.

Related literature. The macroeconomics literature on informational frictions has a long history, a revived present, and—hopefully—a promising future.³ Among this literature, most influential in our thinking have been Morris and Shin (2002), Woodford (2003a), and Angeletos and Pavan (2007, 2009). Morris and Shin (2002) were the first to highlight the potential implications of asymmetric information, and higher-order beliefs, for settings that feature strategic complementarity. Woodford (2003a) exploited the inertia of higher-order beliefs to generate inertia in the response of prices to nominal shocks in a stylized model of price setting. Finally, Angeletos and Pavan (2007a, 2009) provided a methodology for studying the positive and normative properties of a more general class of games with strategic complementarity and dispersed information.

The framework we use in this paper is a variant of the ones we use in two companion papers for different purposes. In Angeletos and La'O (2009b) we elaborate on the broader insight that the heterogeneity of information introduces a distinct type of uncertainty about economic activity, and show how this can sustain sunspot-like fluctuations in a unique-equilibrium economy. In Angeletos and La'O (2008), on the other hand, we extend the analysis of the present paper by allowing information to get aggregated through certain price and quantity indicators and by introducing nominal frictions; we then explore the implications for optimal fiscal and monetary policy.

Part of the combined contribution of these papers is to show how the equilibrium and efficient allocations of fully micro-founded business-cycle economies can be represented as the Perfect Bayesian equilibrium of a game similar to those considered in the more abstract setting of Morris and Shin (2002) and Angeletos and Pavan (2007a, 2009). This representation is useful, as it facilitates a translation of some of the more abstract insights of this earlier work within the context of business cycles. At the same time, the specific micro-foundations are crucial for understanding both the positive and the normative implications of the particular form of complementarity featured in business-cycle models. Indeed, it is only these micro-foundations that explain either why this complementarity turns out to be irrelevant when information is commonly shared, or why it has none of the welfare implications conjectured by Morris and Shin (2002).

Finally, it is worth noting how our approach differs from those of Beaudry and Portier (2004, 2006), Christiano, Motto, and Rostagno (2006), Gilchrist and Leahy (2002), Jaimovich and Rebelo (2008), and Lorenzoni (2008). These papers consider certain types of expectations-driven, or noise-

³See, e.g., Adam (2007), Amador and Weill (2007, 2008), Amato and Shin (2006), Angeletos and Pavan (2004, 2007a, 2007b, 2009), Bacchetta and Wincoop (2005), Coibion and Gorodnichenko (2008), Collard and Dellas (2005), Hellwig (2002, 2005), Hellwig and Veldkamp (2008), Hellwig and Venkateswara (2008), Klenow and Willis (2007), Lorenzoni (2008, 2009), Luo (2008), Mackowiak and Wiederholt (2008, 2009), Mankiw and Reis (2002, 2006), Morris and Shin (2002, 2006), Moscarini (2004), Nimark (2008), Reis (2006, 2008), Rodina (2008), Sims (2003, 2006), Van Nieuwerburgh and Veldkamp (2006), Veldkamp (2006), Veldkamp and Woolfers (2007), and Woodford (2003a, 2008).

driven, fluctuations. However, these fluctuations originate merely from uncertainty about future technological opportunities, obtain within representative-agent models, and do not rest on the heterogeneity of information. Interestingly, Kydland and Prescott (1982) had also allowed for certain expectational shocks; but they, too, did not consider heterogeneous information. Similarly, there are numerous papers that consider geographical and trading structures similar to the one in our model (e.g., Lucas and Prescott, 1974; Rios-Rull and Prescott, 1992; Alvarez and Shimer, 2008), but once again rule out heterogeneous information about the aggregate economic fundamentals.

To recap, it is this particular type of informational heterogeneity that is the distinctive feature of our approach in either the present paper or the aforementioned companion papers. This also distinguishes our approach from the Mirrless literature, which allows for heterogeneous information about idiosyncratic shocks but also rules out heterogeneous information about aggregate shocks.

2 The model

There is a (unit-measure) continuum of households, or “families”, each consisting of a consumer and a continuum of workers. There is a continuum of “islands”, which define the boundaries of local labor markets as well as the “geography” of information: information is symmetric within an island, but asymmetric across islands. Each island is inhabited by a continuum of firms, which specialize in the production of differentiated commodities. Households are indexed by $h \in H = [0, 1]$; islands by $i \in I = [0, 1]$; firms and commodities by $(i, j) \in I \times J$; and periods by $t \in \{0, 1, 2, \dots\}$.

Each period has two stages. In stage 1, each household sends a worker to each of the islands. Local labor markets then open, workers decide how much labor to supply, firms decide how much labor to demand, and local wages adjust so as to clear the local labor market. At this point, workers and firms in each island have perfect information regarding local productivity, but imperfect information regarding the productivities in other islands. After employment and production choices are sunk, workers return home and the economy transits to stage 2. At this point, all information that was previously dispersed becomes publicly known, and commodity markets open. Quantities are now pre-determined by the exogenous productivities and the endogenous employment choices made during stage 1, but prices adjust so as to clear product markets.

Households. The utility of household h is given by

$$u_i = \sum_{t=0}^{\infty} \beta^t \left[U(C_{h,t}) - \int_I S_{i,t} V(n_{hi,t}) di, \right]$$

with

$$U(C) = \frac{C^{1-\gamma}}{1-\gamma} \quad \text{and} \quad V(n) = \frac{n^{1+\epsilon}}{1+\epsilon}.$$

Here, $\gamma \geq 0$ parametrizes the income elasticity of labor supply,⁴ $\epsilon \geq 0$ parameterizes the Frisch elasticity of labor supply, $n_{hi,t}$ is the labor of the worker who gets located on island i during stage 1 of period t , $S_{h,t}$ is an island-specific shock to the disutility of labor, and $C_{h,t}$ is a composite of all the commodities that the household purchases and consumes during stage 2.

This composite, which also defines the numeraire used for wages and commodity prices, is given by the following nested CES structure:

$$C_{h,t} = \left[\int_I c_{hi,t}^{\frac{\rho-1}{\rho}} di \right]^{\frac{\rho}{\rho-1}}$$

where

$$c_{hi,t} = \left[\int_J c_{hij,t}^{\frac{\eta_{it}-1}{\eta_{it}}} dj \right]^{\frac{\eta_{it}}{\eta_{it}-1}}$$

and where $c_{hij,t}$ is the quantity household h consumes in period t of the commodity produced by firm j on island i . Here, η_{it} is a random variable that determines the period- t elasticity of demand faced by any individual firm within a given island i , while ρ is the elasticity of substitution across different islands. Letting the within-island elasticity η differ from the across-islands elasticity ρ permits us to distinguish the degree of monopoly power (which will be determined by the former) from the strength of trade linkages and the associated degree of strategic complementarity (which will be determined by the latter). In fact, a case of special interest that we will concentrate on for much of our analysis is the limit where monopoly power vanishes ($\eta \rightarrow \infty$) while the strategic complementarity remains non-trivial ($\rho < \infty$); this case nests a canonical, competitive RBC economy. At the same time, letting the within-island elasticity to be finite and random permits us to introduce monopoly power and mark-up shocks, thus facilitating a translation/extension of our results to the New-Keynesian framework.

Households own equal shares of all firms in the economy. The budget constraint of household h is thus given by the following:

$$\int_{I \times J} p_{ij,t} c_{hij,t} d(j,k) + B_{h,t+1} \leq \int_{J \times I} \pi_{ij,t} d(i,j) + \int_I w_{it} n_{hi,t} dk + R_t B_{h,t},$$

where $p_{ij,t}$ is the period- t price of the commodity produced by firm j on island i , $\pi_{ij,t}$ is the period- t profit of that firm, w_{it} is the period- t wage on island i , R_t is the period- t nominal gross rate of return on the riskless bond, and $B_{h,t}$ is the amount of bonds held in period t .

The objective of each household is simply to maximize expected utility subject to the budget and informational constraints faced by its members. Here, one should think of the worker-members

⁴Note that risk aversion and intertemporal substitution play no role in our setting because all idiosyncratic risk is insurable and there is no capital. Therefore, γ only controls the sensitivity of labor supply to income for given wage.

of each family as solving a team problem: they share the same objective (family utility) but have different information sets when making their labor-supply choices. Formally, the household sends off during stage 1 its workers to different islands with bidding instructions on how to supply labor as a function of (i) the information that will be available to them at that stage and (ii) the wage that will prevail in their local labor market. In stage 2, the consumer-member collects all the income that the worker-member has collected and decides how much to consume in each of the commodities and how much to save (or borrow) in the riskless bond.

Asset markets. Asset markets operate in stage 2, along with commodity markets, when all information is commonly shared. This guarantees that asset prices do not convey any information. The sole role of the bond market in the model is then to price the risk-free rate. Moreover, because our economy admits a representative consumer, allowing households to trade risky assets in stage 2 would not affect any of the results.

Firms. The output of firm j on island i during period t is given by

$$q_{ij,t} = A_{i,t}(n_{ij,t})^\theta$$

where $A_{i,t}$ is the productivity in island i , $n_{ij,t}$ is the firm's employment, and $\theta \in (0, 1)$ parameterizes the degree of diminishing returns in production. The firm's realized profit is given by

$$\pi_{ij,t} = p_{ij,t}q_{ij,t} - w_{i,t}n_{ij,t}$$

Finally, the objective of the firm is to maximize its expectation of the representative consumer's valuation of its profit, namely, its expectation of $U'(C_t)\pi_{ij,t}$.

Labor and product markets. Labor markets operate in stage 1, while product markets operate in stage 2. Because labor cannot move across islands, the clearing conditions for labor markets are as follows:

$$\int_J n_{ij,t}dj = \int_H n_{hi,t}dh \quad \forall i$$

On the other hand, because commodities are traded beyond the geographical boundaries of islands, the clearing conditions for the product markets are as follows:

$$\int_H c_{hij,t}dh = q_{ij,t} \quad \forall (i, j)$$

Fundamentals and information. Each island in our economy is subject to three types of shocks: shocks to the technology used by local firms (TFP shocks); shocks to the disutility of labor faced by local workers (taste shocks); and shocks to the elasticity of demand faced by local firms, causing variation in their monopoly power (mark-up shocks). We allow for both aggregate and idiosyncratic components to these shocks.

The aggregate fundamentals of the economy in period t are identified by the joint distribution of the shocks $(A_{it}, S_{it}, \eta_{it})$ in the cross-section of islands.⁵ Let Ψ_t denote this distribution. The standard practice in macroeconomics is to assume that Ψ_t is commonly known in the beginning of period t . In contrast, we consider situations where information about Ψ_t is imperfect and, most importantly, heterogeneous. We thus assume that different islands observe only noisy private (local) signals about Ψ_t in stage 1, when they have to make their decentralized employment and production choices. On the other hand, we assume that Ψ_t becomes common known in stage 2, when agents meet in the centralized commodity and financial markets.

For our main theoretical results we do not need to make any special assumptions about the information that is available to each island. For example, we can impose a Gaussian structure as in Morris and Shin (2002). Alternatively, we could allow some islands to be perfectly informed and others to be imperfectly informed, mimicking the idea in Mankiw and Reis (2002) that only a fraction of the agents update their information sets in any given point of time. To some extent, we could even interpret the noise in these signals as the product of rational inattention, as in Sims (2003) and Woodford (2003a). More generally, we do not expect the details of the origins of noise to be crucial for our positive results.

We thus start by allowing a rather arbitrary information structure, as in the more abstract work of Angeletos and Pavan (2009). First, we let ω_t denote the “type” of an island during period t . This variable encodes all the information available to an island about the local shocks as well as about the cross-sectional distribution of shocks and information in the economy. Next, we let Ω_t denote the distribution of ω_t in the cross-section of islands. This variable identifies the aggregate state of the economy during period t ; note that the aggregate state now includes not only the cross-sectional distribution Ψ_t of the shocks but also the cross-sectional distributions of the information (signals). Finally, we let \mathcal{S}_ω denote the set of possible types for each island, \mathcal{S}_Ω the set of probability distributions over \mathcal{S}_ω , and $\mathcal{P}(\cdot|\cdot)$ a probability measure over \mathcal{S}_Ω .⁶

We then formalize the information structure as follows. In the beginning of period t , and conditional on Ω_{t-1} , Nature draws a distribution $\Omega_t \in \mathcal{S}_\Omega$ using the measure $\mathcal{P}(\Omega_t|\Omega_{t-1})$.⁷ Nature then uses Ω_t to make independent draws of $\omega_t \in \mathcal{S}_\omega$, one for each island. In the beginning of period t , before they make their current-period employment and production choices, agents in any given

⁵In special cases (as with Assumption 1 later on), this distribution might be conveniently parameterized by the mean values of the shocks; but in general the aggregate fundamentals are identified by the entire distribution.

⁶To avoid getting distracted by purely technical issues, our proofs treat \mathcal{S}_ω and \mathcal{S}_Ω as if they were finite sets. However, none of our results hinges on this restriction.

⁷Note that we have imposed that the aggregate state Ω_t follows a Markov process; apart from complicating the notation, nothing changes if we let the aforementioned probability measure depend on all past aggregate states.

island get to see only their own ω_t ; in general, this informs them perfectly about their local shocks, but only imperfectly about the underlying aggregate state Ω_t . In the end of the period, however, Ω_t becomes commonly known (ensuring that Ψ_t also becomes commonly known).

To recap, the key informational friction in our model is that agents face uncertainty about the underlying aggregate state Ω_t . Whether they face uncertainty about their own local shocks is immaterial for the type of effects we analyze in this paper. Merely for convenience, then, we assume that the agents of an island learn their own local shocks in stage 1. We can thus express the shocks as functions of ω_t : we denote with $A(\omega_t)$ the local productivity shock, with $S(\omega_t)$ the local taste shock, and with $\eta(\omega_t)$ the local mark-up shock.

3 Equilibrium

In this section we characterize the equilibrium by providing a game-theoretic representation that turns out to be instrumental for our subsequent analysis.

3.1 Definition

Because each family sends workers to every island and receives profits from every firm in the economy, each family's income is fully diversified during stage 2. This guarantees that our model admits a representative consumer and that no trading takes place in the financial market. To simplify the exposition, we thus set $B_t = 0$ and abstract from the financial market. Furthermore, because of the symmetry of preferences, technologies and information within each island, it is without any loss of generality to impose symmetry in the choices of workers and firms within each island. Finally, because of the absence of capital and the Markov restriction on the aggregate state, Ω_{t-1} summarizes all the payoff-relevant public information as of the beginning of period t . Recall then that the additional information that becomes available to an island in stage 1 is only ω_t . As a result, the local levels of labor supply, labor demand, wage, and output can all depend on Ω_{t-1} and ω_t , but not the current aggregate state Ω_t . On the other hand, the commodity prices in stage 2, and all aggregate outcomes, do depend on Ω_t . We thus define an equilibrium as follows.

Definition 1. *An equilibrium consists of an employment strategy $n : \mathcal{S}_\omega \times \mathcal{S}_\Omega \rightarrow \mathbb{R}_+$ a production strategy $q : \mathcal{S}_\omega \times \mathcal{S}_\Omega \rightarrow \mathbb{R}_+$, a wage function $w : \mathcal{S}_\omega \times \mathcal{S}_\Omega \rightarrow \mathbb{R}_+$, an aggregate output function $Q : \mathcal{S}_\Omega^2 \rightarrow \mathbb{R}_+$, an aggregate employment function $N : \mathcal{S}_\Omega^2 \rightarrow \mathbb{R}_+$, a price function $p : \mathcal{S}_\omega \times \mathcal{S}_\Omega^2 \rightarrow \mathbb{R}_+$, and a consumption strategy $c : \mathbb{R}_+^3 \rightarrow \mathbb{R}_+$, such that the following are true:*

(i) The price function is normalized so that

$$P(\Omega_t, \Omega_{t-1}) \equiv \left[\int p(\omega, \Omega_t, \Omega_{t-1})^{1-\rho} d\Omega_t(\omega) \right]^{\frac{1}{1-\rho}} = 1$$

for all (Ω_t, Ω_{t-1}) .

(ii) The quantity $c(p, p', Q)$ is the representative consumer's optimal demand for any commodity whose price is p when the price of all other commodities from the same island is p' and the aggregate output (income) is Q .

(iii) When the current aggregate state is Ω_t and the past aggregate state is Ω_{t-1} , the price that clears the market for the product of the typical firm from island ω_t is $p(\omega_t, \Omega_t, \Omega_{t-1})$; the employment and output levels of that firm are, respectively, $n(\omega_t, \Omega_{t-1})$ and $q(\omega_t, \Omega_{t-1})$, with $q(\omega_t, \Omega_{t-1}) = A(\omega_t)n(\omega_t, \Omega_{t-1})^\theta$; and the aggregate output and employment indices are, respectively,

$$Q(\Omega_t, \Omega_{t-1}) = \left\{ \int q(\omega, \Omega_{t-1})^{\frac{\rho-1}{\rho}} d\Omega_t(\omega) \right\}^{\frac{\rho}{\rho-1}} \quad \text{and} \quad N(\Omega_t, \Omega_{t-1}) = \int n(\omega, \Omega_{t-1}) d\Omega_t(\omega).$$

(iv) The quantities $n(\omega_t, \Omega_{t-1})$ and $q(\omega_t, \Omega_{t-1})$ are optimal from the perspective of the typical firm in island ω_t , taking into account that firms in other islands are behaving according to the same strategies, that the local wage is given by $w(\omega_t, \Omega_{t-1})$, that prices will be determined in stage 2 so as to clear all product markets, that the representative consumer will behave according to consumption strategy c , and that aggregate income will be given by $Q(\Omega_t, \Omega_{t-1})$.

(v) The local wage $w(\omega_t, \Omega_{t-1})$ is such that the quantity $n(\omega_t, \Omega_{t-1})$ is also the optimal labor supply of the typical worker in an island of type ω_t .

Note that condition (i) simply means that the numeraire for our economy is the CES composite defined when we introduced preferences. The rest of the conditions then represent a hybrid of a Walrasian equilibrium for the complete-information exchange economy that obtains in stage 2, once production choices are fixed, and a subgame-perfect equilibrium for the incomplete-information game played among different islands in stage 1.

Let us expand on what we mean by this. When firms in an island decide how much labor to employ and how much to produce during stage 1, they face uncertainty about the prices at which they will sell their product during stage 2 and hence they face uncertainty about the marginal return to labor. Similarly, when workers in an island decide how much labor to supply, they face uncertainty about the real income their household will have in stage 2 and hence face uncertainty about the marginal value of the wealth that they can generate by working more. But then note that firms and workers in each island can anticipate that the prices that clear the commodity markets and the realized level of real income are, in equilibrium, determined by the level of employment

and production in other islands. This suggests that we can solve for the general equilibrium of the economy by reducing it to a certain game, where the incentives of firms and workers in an island depend on their expectations of the choices of firms and workers in other islands. We implement this solution strategy in the following.

Remark. To simplify notation, we often use q_{it} as a short-cut for $q(\omega_t, \Omega_{t-1})$, Q_t as a short-cut for $Q(\Omega_t, \Omega_{t-1})$, \mathbb{E}_{it} as a short-cut for $\mathbb{E}[\cdot | \omega_t, \Omega_{t-1}]$, and so on; also, we drop the indices h and j , because we know that allocations are identical across households, or across firms within an island.

3.2 Characterization

Towards solving for the equilibrium, consider first how the economy behaves in stage 2. The optimal demand of the representative consumer for a commodity from island i whose price is p_{it} when the price of other commodities in the same island is p'_{it} is given by the following:

$$c_{it} = \left(\frac{p_{it}}{p'_{it}} \right)^{-\eta_{it}} \left(\frac{p'_{it}}{P_t} \right)^{-\rho} C_t,$$

where $P_t = 1$ by our choice of numeraire.⁸ In equilibrium, $C_t = Q_t$. It follows that the equilibrium consumption strategy is given by $c(p, p', Q) = p^{-\eta} (p')^{\eta-\rho} Q$. Equivalently, the inverse demand function faced by a firm during period t is

$$p_{it} = (p'_{it})^{1-\frac{\rho}{\eta_{it}}} q_{it}^{-\frac{1}{\eta_{it}}} Q_t^{\frac{1}{\eta_{it}}} \quad (1)$$

Consider now stage 1. Given that the marginal value of nominal income for the representative household is $U'(C_t)$ and that $C_t = Q_t$ in equilibrium, the objective of the firm is simply

$$\mathbb{E}_{it} [U'(Q_t) (p_{it} q_{it} - w_{it} n_{it})].$$

Using (1), we conclude the typical firm on island ω_t maximizes the following objective:

$$\mathbb{E}_{it} \left[U'(Q_t) \left((p'_{it})^{1-\frac{\rho}{\eta_{it}}} Q_t^{\frac{1}{\eta_{it}}} q_{it}^{1-\frac{1}{\eta_{it}}} - w_{it} n_{it} \right) \right], \quad (2)$$

where $q_{it} = A_{it} n_{it}^\theta$. As long as $1 > (1 - \frac{1}{\eta_t})\theta > 0$ (which we assume to be always the case), the above objective is a strictly concave function of n_t , which guarantees that the solution to the firm's problem is unique and that the corresponding first-order condition is both necessary and sufficient. This condition is simply given by equating the expected marginal cost and revenue of

⁸To understand this condition, note that $c'_{it} = \left(\frac{p'_{it}}{P_t} \right)^{-\rho} C_t$ is the demand for the basket of commodities produced by a particular island; the demand for the commodity of a particular firm in that islands is then $c_{it} = \left(\frac{p_{it}}{p'_{it}} \right)^{-\eta_{it}} c'_{it}$.

labor, evaluated under local expectation of the equilibrium pricing kernel:

$$\mathbb{E}_{it} [U'(Q_{it})] w_{it} = \left(\frac{\eta_{it} - 1}{\eta_{it}} \right) \mathbb{E}_{it} \left[U'(Q_t) (p'_{it})^{1 - \frac{\rho}{\eta_{it}}} \left(\frac{Q_t}{q_{it}} \right)^{\frac{1}{\eta_{it}}} \right] (\theta A_{it} n_{it}^{\theta-1}). \quad (3)$$

Next, note that, since all firms within an island set the same price in equilibrium, it must be that $p'_{it} = p_{it}$. Along with (1), this gives

$$p'_{it} = p_{it} = \left(\frac{q_{it}}{Q_t} \right)^{-\frac{1}{\rho}}. \quad (4)$$

This simply states that the equilibrium price of the typical commodity of an island relative to the numeraire is equal to the MRS between that commodity and the numeraire. Finally, note that the optimal labor supply of the typical worker on island i is given by equating the local wage with the MRS between the numeraire and leisure:

$$w_{it} = \frac{S_{it} n_{it}^{\epsilon}}{\mathbb{E}_{it} [U'(Q_t)]} \quad (5)$$

Conditions (4) and (5) give the equilibrium prices and wages as functions of the equilibrium allocation. Using these conditions into condition (3), we conclude that the equilibrium allocation is pinned down by the following condition:

$$S_{it} n_{it}^{\epsilon} = \left(\frac{\eta_{it} - 1}{\eta_{it}} \right) \mathbb{E}_{it} \left[U'(Q_t) \left(\frac{q_{it}}{Q_t} \right)^{-\frac{1}{\rho}} \right] (\theta A_{it} n_{it}^{\theta-1}). \quad (6)$$

This condition has a simple interpretation: it equates the private cost and benefit of effort in each island. To see this, note that the left-hand side is simply the marginal disutility of an extra unit of labor in island i ; as for the right-hand side, $\frac{\eta_{it}-1}{\eta_{it}}$ is the reciprocal of the local monopolistic mark-up, $U'(Q_t) \left(\frac{q_{it}}{Q_t} \right)^{-\frac{1}{\rho}}$ is the marginal utility of an extra unit of the typical local commodity, and $\theta A_{it} n_{it}^{\theta-1}$ is the corresponding marginal product of labor.

Note that condition (6) expresses the equilibrium levels of local employment n_{it} and local output q_{it} in relation to the local shocks and the local expectations of aggregate output Q_t . Using the production function, $q_{it} = A_{it} n_{it}^{\theta}$, to eliminate n_{it} in this condition, and reverting to the more precise notation of Definition 1 (i.e., replacing q_{it} with $q(\omega_t, \Omega_{t-1})$, Q_t with $Q(\Omega_t, \Omega_{t-1})$, A_{it} with $A(\omega_t)$, and so on), we reach the following result.

Proposition 1. *Let*

$$f(\omega) \equiv \log \left\{ \theta^{\frac{\theta}{1-\theta+\epsilon+\gamma\theta}} \left(\frac{\eta(\omega) - 1}{\eta(\omega)} \right)^{\frac{\theta}{1-\theta+\epsilon+\gamma\theta}} S(\omega)^{-\frac{\theta}{1-\theta+\epsilon+\gamma\theta}} A(\omega)^{\frac{1+\epsilon}{1-\theta+\epsilon+\gamma\theta}} \right\}$$

be a composite of all the local shocks hitting an island of type ω and define the coefficient

$$\alpha \equiv \frac{\frac{1}{\rho} - \gamma}{\frac{1}{\rho} + \frac{1-\theta+\epsilon}{\theta}} < 1$$

The equilibrium levels of local and aggregate output are the solution to the following fixed-point problem:

$$\log q(\omega_t, \Omega_{t-1}) = (1 - \alpha) f(\omega_t) + \alpha \log \left\{ \mathbb{E} \left[Q(\Omega_t, \Omega_{t-1})^{\frac{1}{\rho} - \gamma} \middle| \omega_t, \Omega_{t-1} \right]^{\frac{1}{\frac{1}{\rho} - \gamma}} \right\} \quad \forall (\omega_t, \Omega_{t-1}) \quad (7)$$

$$Q(\Omega_t, \Omega_{t-1}) = \left[\int q(\omega, \Omega_{t-1})^{\frac{\rho-1}{\rho}} d\Omega_t(\omega) \right]^{\frac{\rho}{\rho-1}} \quad \forall (\Omega_t, \Omega_{t-1}). \quad (8)$$

This result establishes that the general equilibrium of our economy reduces to a simple fixed-point relation between local and aggregate output. In so doing, it offers a game-theoretic representation of our economy, similar to the one established in Angeletos and La'O (2009b) for a variant economy with capital. To see this, consider a game with a large number of players, each choosing an action in \mathbb{R}_+ . Identify a “player” in this game with an island in our economy and interpret the level of output of that island as the “action” of the corresponding player. Next, identify the “types” of these players with ω_t , which encodes the local shocks and local information sets in our economy. Finally, let their “best responses” be given by condition (7). It is then evident that the Perfect Bayesian equilibrium of this game identifies the general equilibrium of our economy.

Note then that the variable $f(\omega_t)$ conveniently summarizes all the local economic fundamentals, while the coefficient α identifies the degree of strategic complementarity in our economy. To see this more clearly, consider a log-linear approximation to conditions (7) and (8):

$$\log q_{it} = \text{const} + (1 - \alpha) f_{it} + \alpha \mathbb{E}_{it} [\log Q_t], \quad (9)$$

$$\log Q_t = \text{const} + \int \log q_{it} di, \quad (10)$$

where *const* capture second- and higher-order terms.⁹ It is then evident that the coefficient α identifies the slope of an island’s best response to the activity of other islands—which is the standard definition of the degree of strategic complementarity.

Finally, note that Proposition 1 holds no matter the information structure. This is important. While much of the recent literature has focused on specific formalizations of the information structure (e.g. Mankiw and Reis, 2002; Sims, 2003; Woodford, 2003a), our result indicates that the

⁹In general, these second- and higher-order terms may depend on the underlying state and the above is only an approximation. However, when the underlying shocks and signals are jointly log-normal with fixed second moments (as imposed by Assumption 1 in the next section), these terms are invariant, the approximation error vanishes, and conditions (9) and (10) are exact.

information structure typically matters *only* by pinning down the agents' forecasts of economic activity. We would thus invite future researchers to pay more attention on the theoretical and empirical properties of these forecasts as opposed to the details of the information structure.

3.3 Trade links and strategic complementarity

As evident from Proposition 1, the degree of complementarity, α , is a monotone function of the elasticity of substitution across the commodities of different islands, ρ . In what follows, we adopt the convention that variation in α represents variation in ρ for given other parameters. We also interpret α as a measure of the strength of trade linkages in our economy. These choices are motivated by the following observations. First, if we consider a variant of our model where each household lives and works only in one island and consumes only the products of that island, then Proposition 1 holds with $\alpha = 0$; in this sense, it is precisely the trade linkages across different islands that introduces strategic interdependence ($\alpha \neq 0$). Second, while α depends, not only on ρ , but also on ϵ, γ , and θ , these other parameters affect the composite shock f and matter for equilibrium allocations whether islands (agents) are linked or not; in contrast, ρ affects only α . For these reasons, we henceforth use the notions of strategic complementarity, elasticity of substitution across islands, and strength of trade linkages, as synonymous to one another. However, we also note that strong complementarity in our model does not strictly require low ρ : if the wealth effect of labor supply is small ($\gamma \rightarrow 0$), the Frisch elasticity is high ($\epsilon \rightarrow 0$), and production is nearly linear ($\theta \rightarrow 1$), then the degree of complementarity is high ($\alpha \rightarrow 1$) no matter what ρ is.

The insight that trade introduces a form of strategic complementarity even in neoclassical, perfectly-competitively settings is likely to extend well beyond the boundaries of the model we have considered here or the variant in Angeletos and La'O (2009b). We believe that this insight has been under-appreciated in prior work on business cycles for two reasons. First, the two welfare theorems have taught us that it rarely helps, and it can often be misleading, to think of Walrasian settings as games. And second, the type of strategic complementarity we highlight here is simply irrelevant for the business cycle when information is commonly shared.

To understand what we mean by the last point, consider the response of the economy to a *symmetric* aggregate shock (i.e., a shock that keeps the level of heterogeneity invariant). Formally, let \bar{f}_t denote the cross-sectional average of the composite fundamental f_{it} and consider any shock that varies the average fundamental, \bar{f}_t , without varying the cross-sectional distribution of the idiosyncratic components of the fundamentals, $\xi_{it} \equiv f_{it} - \bar{f}_t$. When all information is commonly

shared, aggregate output is also commonly known in equilibrium. Condition (7) then reduces to

$$\log q_{it} = (1 - \alpha)(\bar{f}_t + \xi_{it}) + \alpha \log Q_t. \quad (11)$$

It is then immediate that the entire cross-sectional distribution of $\log q_{it}$ moves one-to-one with \bar{f}_t , which establishes the following.

Proposition 2. *Suppose that information is commonly shared and that the level of heterogeneity is invariant. Then the equilibrium levels of aggregate output is given by*

$$\log Q_t = \text{const} + \bar{f}_t.$$

Recall that, by its definition, the composite shock depends on ϵ and γ but not on ρ . It is then evident that the response of the economy to the underlying aggregate productivity, taste, or mark-up shocks is independent of ρ . In this sense, the business cycle is indeed independent of the degree of strategic complementarity that is induced by trade.

The intuition behind this result is further explained in Angeletos and La'O (2009b). The key is that the strength of trade linkages matters only for how much agents care about forecasting the level of economic activity *relatively* to forecasting the underlying economic fundamentals. But when information is symmetric (commonly shared), any uncertainty the agents face about the level of economic activity reduces to the one that they face about the underlying economic fundamentals, which renders the degree of strategic complementarity irrelevant. In contrast, when information is asymmetric (dispersed), agents can face additional uncertainty about the level of economic activity, beyond the one they face about the fundamentals. The strength of trade linkages then dictates precisely the impact on equilibrium outcome of this residual uncertainty about economic activity.

This is important. It is precisely the aforementioned property that makes dispersed information distinct from uncertainty about the fundamentals—for it is only the *heterogeneity* of information that breaks the coincidence of forecasts of economic activity with the forecasts of the underlying fundamentals when the equilibrium is unique. We further elaborate on this point in Angeletos and La'O (2009b), showing how dispersed information can open the door to a certain type of sunspot-like fluctuations. We refer the reader to that paper for a more thorough discussion of this important, broader insight. In what follows, we concentrate on how this broader insight helps understand why the combination of dispersed information with the aforementioned type of complementarity can have a significant impact on the positive properties of the RBC paradigm.

3.4 Relation to complementarity in New-Keynesian models

The familiar condition that characterizes optimal target prices in the New-Keynesian paradigm (e.g., Woodford, 2003b) looks like the following:

$$p_{i,t} = (1 - \xi)\mathcal{Y}_t + \xi p_t + z_{i,t}, \quad (12)$$

where $p_{i,t}$ is the target price of a firm (in logs), \mathcal{Y}_t is nominal GDP, p_t is the aggregate price level, $z_{i,t}$ captures idiosyncratic productivity or demand shocks, and ξ is a coefficient that is interpreted as the degree of strategic complementarity in pricing decisions. If we compare the above condition with condition (9) in our model, the resemblance is striking. The only noticeable difference seems to be that the relevant choice variable is a price in the New-Keynesian model, while it is a quantity in our model. However, there are some crucial differences behind this resemblance.

First, condition (12) does not *alone* pin down the equilibrium. Rather, it must be combined with other conditions regarding the determination of \mathcal{Y}_t , the nominal GDP level. In contrast, condition (9) offers a complete, self-contained, representation of the equilibrium in our model.

Second, the endogeneity of \mathcal{Y}_t undermines the meaning of condition (12). For example, letting y_t denote real GDP and using $\mathcal{Y}_t = p_t + y_t$, condition (12) can also be restated as $p_{i,t} = p_t + (1 - \xi)y_t + z_{i,t}$; but then the degree of complementarity appears to be 1, not ξ . In fact, this alternative representation is more informative when money is neutral, because y_t is then exogenous to nominal factors and this condition determines only relative prices. But even when money is non-neutral, ξ fails to identify the degree of complementarity in pricing decisions simply because nominal GDP is far from exogenous—at the very least because monetary policy responds to variation in p_t and y_t . Once this endogeneity is incorporated, the complementarity in pricing decisions is different from ξ and becomes sensitive to policy parameters. In contrast, in our model the degree of strategic complementarity is pinned down only by preferences and technologies, and is completely invariant to monetary policy.

Third, the comparative statics of the complementarity in our model (α) with respect to deeper preference and technology parameters are different from those of its New-Keynesian counterpart (ξ). In particular, note that α decreases with ρ (the elasticity of substitution across different goods), decreases with ϵ (the inverse of the Firsch elasticity of labor supply), and increases with θ (the degree of diminishing returns to labor). Hence, what contributes to strong complementarity in our model is low substitutability in the commodity side, so that trade is crucial, along with high substitutability in the labor and production side, as in Hansen (1986) and King and Rebelo (2000). As one of our discussants highlighted, the opposite comparative statics hold for ξ in the New-Keynesian paradigm. Nevertheless, it is important to bear in mind that our notion of complementarity may have little to do with either the degree of monopoly power or the price elasticities of individual demands. In

our model, that latter are pinned down by η (the within-island elasticity of substitution), while the degree of strategic complementarity is pinned down by ρ (the across-island elasticity).

Last, but not least, the complementarity highlighted in the New-Keynesian framework would vanish if firms were setting real (indexed) prices. In this sense, the New-Keynesian complementarity in is a nominal phenomenon, whereas ours is a real phenomenon.

4 Dispersed information and the business cycle

In this section we seek to illustrate how the introduction of dispersed information can impact the positive properties of the RBC paradigm. To facilitate this task, we impose a Gaussian specification on the shocks and the information structure, similar to the one in Morris and Shin (2002), Woodford (2003a), Angeletos and Pavan (2007), and many others.

Assumption 1. *The shocks and the available information satisfy the following properties:*

(i) *The aggregate shock \bar{f}_t follows a Gaussian AR(1) or random walk process:*

$$\bar{f}_t = \psi \bar{f}_{t-1} + \nu_t,$$

where ψ parameterizes the persistence of the composite shock and ν_t is a Normal innovation, with mean 0 and variance $\sigma_\nu^2 \equiv 1/\kappa_f$, i.i.d. over time.

(ii) *The local shock f_t is given by*

$$f_t = \bar{f}_{it} + \xi_{it},$$

where ξ_{it} is a purely idiosyncratic shock, Normally distributed with mean zero and variance σ_ξ^2 , orthogonal to \bar{f}_t , and i.i.d. across islands.

(iii) *The private information of an island about the aggregate shock \bar{f}_t is summarized in a Gaussian sufficient statistic x_{it} such that*

$$x_{it} = \bar{f}_t + \varsigma_{it},$$

where ς_{it} is noise, Normally distributed with mean zero and variance $\sigma_\varsigma^2 \equiv 1/\kappa_x$, orthogonal to both \bar{f}_t and ξ_{it} , and i.i.d. across islands.¹⁰

(iv) *The public information about the aggregate shock \bar{f}_t is summarized in a Gaussian sufficient statistic y_t such that*

$$y_t = \bar{f}_t + \varepsilon_t,$$

where ε_t is noise, Normally distributed with mean zero and variance $\sigma_\varepsilon^2 \equiv 1/\kappa_y$, and orthogonal to all other variables.

¹⁰Note that the local fundamental f_{it} is itself a private signal of \bar{f}_t . However, by the fact that we define x_{it} as a sufficient statistic of *all* the local private information, the informational content of f_{it} is already included in x_{it} .

This specification imposes a certain correlation in the underlying productivity, taste and mark-up shocks: for the composite shock f_{it} to follow a univariate process as above, it must be that all the three type of shocks are moved by a single underlying factor. However, this is only for expositional simplicity. We can easily extend our results to a situation where each of the shocks follows an independent Gaussian process, or consider a more general correlation structure among the shocks.

4.1 Closed-form solution

Under Assumption 1, we can identify ω_t with the vector (f_t, x_t, y_t) . Because Ω_t is then a joint normal distribution with mean $(\bar{f}_t, \bar{f}_t, y_t)$ and an invariant variance-autocovariance matrix, we can also reduce the aggregate state variable from Ω_t to the more convenient vector (\bar{f}_t, y_t) . Next, we can guess and verify that there is always an equilibrium in which $\log q_{it}$ is linear in $(\bar{f}_{t-1}, f_{it}, x_{it}, y_t)$ and $\log Q_t$ is linear in $(\bar{f}_{t-1}, \bar{f}_t, y_t)$. We then find the coefficients of these linear functions by the familiar method of undetermined coefficients. Finally, we can use an independent argument to rule out any other equilibrium. We thereby reach the following result.

Proposition 3. *Under Assumption 1, the equilibrium level of local output is given by*

$$\log q_{it} = \text{const} + \varphi_{-1}\bar{f}_{t-1} + \varphi_f f_{it} + \varphi_x x_{it} + \varphi_y y_t, \quad (13)$$

where the coefficients $(\varphi_{-1}, \varphi_f, \varphi_x, \varphi_y)$ are given by

$$\begin{aligned} \varphi_{-1} &= \left\{ \frac{\kappa_f}{(1-\alpha)\kappa_x + \kappa_y + \kappa_f} \right\} \alpha \psi & \varphi_f &= (1-\alpha) \\ \varphi_x &= \left\{ \frac{(1-\alpha)\kappa_x}{(1-\alpha)\kappa_x + \kappa_y + \kappa_f} \right\} \alpha & \varphi_y &= \left\{ \frac{\kappa_y}{(1-\alpha)\kappa_x + \kappa_y + \kappa_f} \right\} \alpha \end{aligned} \quad (14)$$

This result gives a closed-form solution of the equilibrium level of output in each island as a log-linear function of the past aggregate fundamental \bar{f}_{t-1} , the current local fundamental f_{it} , the local (private) signal x_{it} , and the public signal y_t . Note then that the equilibrium level of output is necessarily an increasing function of the local fundamental f_{it} : $\varphi_f > 0$ necessarily. To interpret this sign, note that higher f means a higher productivity, a lower disutility of labor, or a lower monopolistic distortion. But whether and how local output depends on \bar{f}_{t-1} , x_{it} and y_t is determined by the degree of strategic complementarity α .

To understand this, note that local output depends on these variables only because these variables contain information about the current aggregate shocks and, in so doing, help agents forecast the aggregate level of output. But when $\alpha = 0$, the demand- and supply side effects that we discussed earlier perfectly offset each other, so that at the end economic decisions are not interdependent: local incentives depend only the local fundamentals and not on expectations of aggregate activity.

It follows that the dependence of local output to \bar{f}_{t-1} , x_{it} and y_t vanishes when $\alpha = 0$. On the other hand, if $\alpha \neq 0$, local output depends on \bar{f}_{t-1} , x_{it} and y_t because, and only because, these variables help predict aggregate output. In particular, when economic decisions are strategic complements ($\alpha > 0$), the equilibrium level of output in each island responds positively to expectations of aggregate output; in this case, the coefficients φ_{-1} , φ_x , and φ_y are all positive. When instead economic decisions are strategic substitutes ($\alpha < 0$), the equilibrium level of output in each island responds negatively to expectations of aggregate output; in this case, the coefficients φ_{-1} , φ_x , and φ_y are all negative. As mentioned earlier, we view the case in which $\alpha > 0$, and hence in which economic activity responds positively to good news about aggregate fundamentals, as the empirically most relevant scenario. For this reason, our subsequent discussion will focus on this case; however, our results apply more generally.

4.2 Remark on interpretation of noise and comparative statics

Before we proceed, we would like to emphasize that one should not give a narrow interpretation to the signal y_t , or its noise ε_t . This signal is not meant to capture only purely public information; rather, it is a convenient modeling device for introducing correlated errors in beliefs of aggregate fundamentals. Indeed, the results we document below can easily be re-casted with a more general information structure, one that allows agents to observe multiple private signals and introduce imperfect cross-sectional correlation in the errors of these private signals; the origin of noise, then, is not only the public signal, but also the correlated errors in the private signals of the agents. We invite the reader to keep this more general interpretation of what “noise” stands for in our model: it is a acronym for all sources of correlated errors in expectations of the fundamentals.¹¹

Similarly, we would like to warn the reader not to focus on the comparative statics of the equilibrium with respect to the precisions of private and public information, κ_x and κ_y , as one of our discussants suggests. These comparative statics fail to isolate the distinct impact of the *heterogeneity* of information, simply because they confound a change in the heterogeneity of information with a change in the overall precision of information.¹² Furthermore, if we had allowed for multiple private signals with correlated errors, it would be unclear whether an increase in the precision of a certain

¹¹In fact, one could go further and interpret “noise” as a certain type of sentiment shocks, namely shocks that do not move at all the agents’ beliefs about the fundamentals and nevertheless move equilibrium outcomes. With a unique-equilibrium model as ours, such shocks cannot exist when information is commonly shared; but emerge robustly once information is dispersed. See Angeletos and La’O (2009b).

¹²For example, an increase in κ_x would increase the heterogeneity of information, but would also increase the overall precision of information; and while the former effect would tend to amplify the volatility effects we have documented here, the latter effect would work in the opposite direction.

signal raises or reduces the heterogeneity of information. With this in mind, in what follows we focus on the comparative statics with respect to α . These comparative statics best isolate the distinct role of dispersed information, simply because the degree of complementarity matters for aggregate fluctuations in our model *only* by regulating the impact of the heterogeneity of information.¹³

4.3 Macroeconomic responses to fundamentals and noise

We now study how the dispersion of information and the strength of trade linkages affect aggregate fluctuations. Towards this goal, we aggregate condition (13) and use the fact that $\bar{f}_t = \psi \bar{f}_{t-1} + \nu_t$ to obtain the following characterization of aggregate output.

Corollary 1. *Under Assumption 1, the equilibrium level of aggregate output is given by*

$$\log Q_t = \text{const} + \psi \bar{f}_{t-1} + \varphi_\nu \nu_t + \varphi_\varepsilon \varepsilon_t, \quad (15)$$

where

$$\varphi_\nu \equiv \varphi_f + \varphi_x + \varphi_y = 1 - \frac{\alpha \kappa_f}{(1 - \alpha) \kappa_x + \kappa_y + \kappa_f} \quad \text{and} \quad \varphi_\varepsilon \equiv \varphi_y = \frac{\alpha \kappa_y}{(1 - \alpha) \kappa_x + \kappa_y + \kappa_f}, \quad (16)$$

and where $\nu_t = \bar{f}_t - \psi \bar{f}_{t-1}$ is the innovation in the fundamentals, ψ is the persistence in the fundamentals, $\varepsilon_t = y_t - \bar{f}_t$ is the aggregate noise.

Condition (15) gives the equilibrium level of aggregate output as a log-linear function of the past aggregate fundamentals, \bar{f}_{t-1} , the current innovation in the fundamentals, ν_t , and the current noise, ε_t . Consider the impact effect of an innovation in fundamentals. This effect is measured by the coefficient φ_ν . Because the latter is a decreasing function of the precisions κ_x and κ_y , we have that the impact effect of an innovation in fundamentals decreases with the level of noise. This is essentially the same insight as the one that drives the real effects of monetary shocks in both the older macro models with informational frictions (e.g., Lucas, 1972; Barro, 1976) and their recent descendants (e.g., Mankiw and Reis, 2002): the less informed economic agents are about the underlying shocks, the less they respond to these shocks. Clearly, this is true no matter whether agents interact with one another—it is true even in a single-agent decision problem.

¹³Angeletos and Pavan (2007a) propose that a good measure of the “commonality” of information (an inverse measure of the heterogeneity of information) is the cross-sectional correlation of the errors in the agents’ forecasts of the fundamentals: holding constant the variance of these forecast errors, an increase in the correlation implies that agents can better forecast one another’s actions, even though they cannot better forecast the fundamentals. Following this alternative route would deliver similar insights as the ones we document here.

More interestingly, we find that φ_ν is a decreasing function of α . That is, the more economic agents care about aggregate economic activity, the weaker the response of the economy to innovations in the underlying fundamentals. At the same time, we find that φ_ε is an increasing function of α . That is, the more economic agents care about aggregate economic activity, the stronger the equilibrium impact of noise. These properties originate from the interaction of strategic complementarity with dispersed information. Indeed, if the underlying shock was common knowledge (which here can be nested by taking the limit as the public signal becomes infinitely precise, $\kappa_y \rightarrow \infty$), then both φ_ν and φ_ε would cease to depend on α . But as long as information is dispersed, a higher α reduces φ_ν and raises φ_ε . This highlights how strategic complementarity becomes crucial for the business cycle once information is dispersed.

Corollary 2. *When information is dispersed, and only then, stronger complementarity dampens the impact of innovations in the fundamentals on equilibrium output and employment, while amplifying the impact of noise.*

The key intuition behind this result is the same as the one in the more abstract work of Morris and Shin (2002) and Angeletos and Pavan (2007a). Public information and past fundamentals (which here determine the prior about the current fundamentals) help forecast the aggregate level of output relatively better than private information. The higher α is, the more the equilibrium level of output in any given island depends on the local forecasts of aggregate output and the less it depends on the local current fundamentals. It follows that a higher α induces the equilibrium output of each island to be more anchored to the past aggregate fundamentals, more sensitive to public information, and less sensitive to private information. The anchoring effect of past aggregate fundamentals explains why aggregate output responds less to any innovation in the fundamentals, while the heightened sensitivity to noisy public information explains why aggregate output responds more to noise. A similar anchoring effect of the common prior underlies the inertia effects in Woodford (2003a), Morris and Shin (2006), and Angeletos and Pavan (2007a), while the heightened sensitivity to public information is the same as the one in Morris and Shin (2002). However, as mentioned before, we favor a more general interpretation of the signal y_t , not as a public signal, but rather as a source of correlated noise in forecasts of economic fundamentals.

As another way to appreciate the aforementioned result, consider following variance-decomposition exercise. Let $\log \hat{Q}_t$ be the projection of $\log Q_t$ on *past* fundamentals. The residual, which is given by $\log \tilde{Q}_t \equiv \log Q_t - \log \hat{Q}_t = \varphi_\nu \nu_t + \varphi_\varepsilon \varepsilon_t$, can be interpreted as the "high-frequency component" of aggregate output. Its total variance is $Var(\log \tilde{Q}_t) = \varphi_\nu^2 \sigma_\nu^2 + \varphi_\varepsilon^2 \sigma_\varepsilon^2$, where $\sigma_\nu^2 (\equiv 1/\kappa_f)$ is the variance of the innovation in the fundamentals and $\sigma_\varepsilon^2 (\equiv 1/\kappa_y)$ is the variance of the noise. The fraction of

the high-frequency variation in output that originates in noise is thus given by the following ratio:¹⁴

$$R_{noise} \equiv \frac{Var(\log \tilde{Q}_t | \nu_t)}{Var(\log \tilde{Q}_t)} = \frac{\varphi_\varepsilon^2 \sigma_\varepsilon^2}{\varphi_\nu^2 \sigma_\nu^2 + \varphi_\varepsilon^2 \sigma_\varepsilon^2}.$$

Since a higher α raises φ_ε and reduces φ_ν , it necessarily raises this fraction: the more agents care about the aggregate level of economic activity, the more the high-frequency volatility in output that is driven by noise.

We can then further highlight the distinct nature of dispersed information by showing that, as long as α is high enough, the contribution of noise to short-run fluctuations can be large even if the level of noise is small. Note that the overall precision of an agent's posterior about the underlying fundamentals is given by $\kappa = \kappa_0 + \kappa_x + \kappa_y$. We can then show the following.

Proposition 4. *When information is dispersed and α is sufficiently high, agents can be arbitrarily well informed about the fundamentals ($\kappa \approx \infty$) and, yet, the high-frequency variation in aggregate output can be driven almost exclusively by noise ($R_{noise} \approx 1$).*

Clearly, this is not possible when information is commonly shared. In that case, the contribution of noise on the business cycle is tightly connected to the precision of information and vanishes as this precision becomes infinite. In contrast, when information is dispersed, the contribution of noise in the business cycle can be high even when the precision of information is arbitrarily high. What makes this possible is the combination of heterogeneous information with a sufficiently strong degree of strategic complementarity induced by trade linkages. Note then how this result also contrasts with our earlier observation that this particular type of strategic complementarity would have been irrelevant for the business cycle had information been commonly shared.

Finally, it is worth noting how the dispersion of information and trade linkages affect the cyclical behavior of aggregate employment. The latter is given by

$$\log N_t = \text{const} + \frac{1}{\bar{\theta}}(\log Q_t - \bar{a}_t),$$

where \bar{a}_t is the aggregate productivity shock (i.e., the cross-sectional average of $\log A_{i,t}$). It is then immediate that the response of employment to an aggregate shock in either tastes or monopoly power is proportional to that of output. The same is true for the response to noise. More interestingly, the response of employment to an aggregate productivity shock may now turn from a positive sign under common information to a negative sign under dispersed information. To see this, let $\beta \equiv \frac{\partial \bar{f}_t}{\partial \bar{a}_t} = \frac{1+\epsilon}{1-\theta+\epsilon+\theta\gamma} > 0$. When information is commonly shared, the sensitivity of output to an innovation to aggregate productivity is simply β , and that of employment is $\frac{1}{\bar{\theta}}(\beta - 1)$. When,

¹⁴This fraction equals 1 minus the R-square of the regression of $\log \tilde{Q}_t$ on the innovation ν_t .

instead, information is dispersed, the corresponding sensitivities are $\varphi_\nu\beta$ for output and $\frac{1}{\theta}(\varphi_\nu\beta - 1)$ for employment, with φ_ν as in (16). Suppose $\beta > 1$, which means that employment responds positively to a productivity shock under common information, as in any plausible calibration of the RBC framework. As noted earlier, φ_ν is necessarily lower than 1 and is decreasing in α . It follows that, when information is dispersed, stronger trade linkages dampen the response of employment and may actually turn it negative.

5 Slow learning

The preceding has focused on a setting where the underlying shocks become common knowledge within a period. Although this permitted a sharp theoretical analysis of the distinct implications of dispersed information, and of its interaction with trade linkages, it makes it hard to map our results to either empirical business cycles or calibrated RBC models. We now seek to illustrate how incorporating slower learning can facilitate a better mapping between our analysis and the data.

Towards this goal, we need to relax the assumption that the aggregate state, Ω_t , becomes publicly revealed at the end of each period, and instead allow for more interesting learning dynamics. Accommodating this possibility in a fully micro-founded way would require that there is no centralized commodity trading: with centralized trading, equilibrium prices are likely to reveal the state. However, allowing for decentralized trading would complicate the analysis by introducing informational externalities and/or by letting the relevant state space explode as in Townsend (1983). We are currently exploring some possibilities along these lines. However, for the current purposes, we opt for tractability and expositional simplicity.

In particular, we assume that firms and workers do not observe the state at the end of each period, nor do they learn about it from observing past aggregate economic outcomes or past prices. Rather, they only keep receiving exogenous signals about the current fundamentals, of the same type as in Assumption 1, and they use these signals to update each period their beliefs about the underlying state. Think of this as follows. Each firm has two managers: one who decides the level of employment and production; and another who sells the product, receives the revenue, and sends the realized profits to the firm's shareholders. The two managers share the same objective—maximize firm valuation—but do not communicate with one another. Moreover, the first manager never receives any signals on economic activity. He only observes the exogenous local private and public signals. Similarly, the consumers, who observe all the prices in the economy, fail to communicate this information to the workers in their respective families. The workers also base their decisions solely on the exogenous signals.

Needless to say, this specification of the learning process is not particularly elegant. However, it would also be naive to take it too literally: the exogenous signals that we allow firms and workers to receive each period are meant to capture more generally the multiple sources of information that these agents may have. To the extent that the underlying shocks do not become common knowledge too fast, more plausible formalizations of the learning process, albeit highly desirable, need not impact the qualitative properties we wish highlight here.¹⁵

Under the aforementioned specification, equilibrium behavior continues to be characterized by the same best-response-like condition as in the baseline model:

$$\log q_{i,t} = (1 - \alpha)f_{i,t} + \alpha\mathbb{E}_{i,t}[\log Q_t], \quad (17)$$

where we have normalized the constant to zero. The only difference is in the information that underlies the expectation operator in this condition. Finally, for concreteness, we henceforth focus on productivity shocks as the only shock to fundamentals: $f_{i,t} = \beta \log A_{i,t}$, with $\beta \equiv \frac{1+\epsilon}{1-\theta+\epsilon+\theta\gamma}$.

The procedure we follow to solve for the equilibrium dynamics is based on Kalman filtering and is similar to the one in Woodford (2003a). We guess and verify that the aggregate state can be summarized in a vector X_t comprised of the aggregate fundamental and aggregate output:

$$X_t \equiv \begin{bmatrix} \bar{f}_t \\ \log Q_t \end{bmatrix}, \quad (18)$$

Firms and workers in any given island never observe the state, but instead receive the following vector of signals each period:

$$z_{it} \equiv \begin{bmatrix} x_{it} \\ y_t \end{bmatrix} = \begin{bmatrix} \bar{f}_t + \varsigma_{it} \\ \bar{f}_t + \varepsilon_t \end{bmatrix} \quad (19)$$

As emphasized before, y_t should not be taken too literally—it is a convenient modeling device for introducing common noise in the agents' forecasts of the state of the economy. Finally, we guess and verify that the state vector X_t follows a simple law of motion:

$$X_t = MX_{t-1} + m_\nu \nu_t + m_\varepsilon \varepsilon_t \quad (20)$$

where M is a 2×2 matrix, while m_ν and m_ε are 2×1 vectors. We then seek to characterize the equilibrium values of M , m_ν , and m_ε .

¹⁵The learning process we assume here is similar to the one in Woodford (2003a). We refer the reader to Amador and Weill (2008), Angeletos and La'O (2008), Angeletos and Pavan (2009), Hellwig (2002), and Lorenzoni (2008) for some alternative formalizations of the learning process. None of these alternative formalizations would crucially affect the positive results we document in this section. However, some of them would have distinct normative implications, to the extent that they make the learning endogenous to the actions of other agents. See Angeletos and La'O (2008) and Angeletos and Pavan (2009) on this issue.

In each period t , firms and workers start with some prior about X_t and use the new signals that they receive in the beginning of period t to update their beliefs about X_t . Local output is then determined. Condition (17) then gives local output as a function of the local belief about X_t . Aggregating across islands, we obtain the aggregate level of output. In equilibrium, the law of motion that aggregate output follows must match the one believed by the firms. Therefore the equilibrium is a fixed point between the law of motion believed by agents and used to form their forecasts of the aggregate state, and the law of motion induced by the optimal output and employment decisions that firms and workers are making following their signal extraction problem. We characterize the fixed point of this problem in the Appendix and use its solution to numerically simulate the impulse responses of output and employment to positive innovations in v_t and ε_t .

For our numerical simulations, we interpret a period as a quarter. Accordingly, we let $\sigma_v = 0.02$ for the standard deviation of the productivity innovation and $\psi = 0.99$ for its persistence. Next, we set $\theta = .60$ and $\epsilon = .5$, which correspond to an income share of labor equal to 60% and a Frisch elasticity of labor supply equal to 2. These parameter values are broadly consistent with the literature. Less standard is our choice of γ . Recall that in our setting there is no capital, implying that labor income is the only source of wealth, the elasticity of intertemporal substitution is irrelevant, and γ only controls the income elasticity of labor supply. We accordingly set $\gamma = .2$ to ensure an empirically plausible income effect on labor supply.¹⁶ Next, we set the standard deviations of the noises as $\sigma_x = \sigma_y = 5\sigma_v$. These values are arbitrary, but they are not implausible: when the period is interpreted as a quarter, the information about the current innovations to fundamentals and/or the current level of economic activity is likely to be very limited. Finally, we do not pick any specific value for α (equivalently, ρ). Rather, we study how the variance decomposition of the high-frequency components of output and employment varies as we vary α from 0 to 1 (keeping in mind that a higher α means stronger trade linkages or, equivalently, a lower ρ).

5.1 Impulse responses to productivity and noise shocks

Figure 1 plots the impulse responses of aggregate output and employment to a positive innovation of productivity, for various degrees of α . (The size of the innovation here, and in all other impulse responses we report, is equal to one standard deviation.) Clearly, if aggregate productivity were common knowledge, then output would follow the same AR(1) process as aggregate productivity itself. This is simply because there is no capital in our model. The same thing happens when information is dispersed but there is no strategic complementarity in output decisions ($\alpha = 0$).

¹⁶Woodford (2003a) uses a similar number for calibrating the New-Keynesian model in the absence of capital.




Figure 1: Impulse responses to a positive innovation in productivity.

This is simply because when $\alpha = 0$ islands are effectively isolated from one another; but as each island knows perfectly its own productivity, the entire economy responds to the aggregate shock as if the aggregate shock had been common knowledge.

In contrast, when information is dispersed but islands are interconnected ($\alpha \neq 0$), employment and output in one island depends crucially on expectations of employment and output in other islands. As a result, even though each island remains perfectly informed about their local fundamentals, each island responds less to the shock than what it would have done had the shock been common knowledge, precisely because each island expects output in other islands to respond less. Note then that the key for the response of each island is not per se whether the island can disentangle an aggregate shock from an idiosyncratic shock. Even if a particular island was perfectly informed about the aggregate shock, as long as $\alpha > 0$ the island will respond less to this shock than under common knowledge if it expects the other island to respond less, presumably because the other island has imperfect information about the shock. Thus, the key for the inertia in the response of aggregate outcomes is the uncertainty islands face about one another's response, not necessarily the uncertainty they themselves face about the aggregate shock.

As evident in Figure 1, the equilibrium inertia is higher the higher the degree of strategic complementarity. This is because of two reasons. First, there is a direct effect: the higher α is, the less the incentive of each island to respond to the underlying shock for any given expectation of the response of other islands. But then there is also an indirect, multiplier-like, effect: as all islands are




Figure 2: Impulse responses to noise.

expected to respond less to the underlying shock, each island finds it optimal to respond even less.

At the same time, the inertia vanishes in the long-run: the long-run response of the economy to the shock is the same as with common knowledge. This seems intuitive: as time passes, agents become better informed about the underlying aggregate shock. However, that's only part of the story. First, note that agents are always perfectly informed about their own fundamentals, so there is no learning in this dimension. Second, recall that agents do not care per se about the aggregate fundamentals, so the fact that they are learning more about them is per se inconsequential. Rather, the key is that agents in each island are revising their forecasts of the output of other islands. What then drives the result that inertia vanishes in the long-run is merely that forecasts of aggregate output eventually converge their common-knowledge counterpart.¹⁷

Finally, a salient property of the dynamic response of employment is that, for sufficiently high α , the short-run impact of a productivity shock on employment turns from positive to negative; this happens for parameters values for which the model would have generate a strong positive response had information been symmetric. We find this striking. The baseline RBC paradigm has long been

¹⁷It may be hard to fully appreciate this point, because how fast output forecasts converge to their common-knowledge counterpart is itself pinned down by the speed of learning about the underlying aggregate productivity shock. However, with richer information structures, one can disentangle the speed of adjustment in output forecasts from the speed of learning about the fundamentals. It is then only the former that matters for the result. See Angeletos and La'O (2009a) for a related example within the context of a Calvo-like monetary model.

criticized for generating a near perfect correlation between employment and output, whereas in the data this correlation is near zero. In our setting, this correlation could be close to zero or even turn negative if α is sufficiently high. Of course, correlations confound the effects of multiple shocks. Some authors in the structural VAR literature have thus sought to show that identified technology shocks lead to a reduction in employment and have then argue that this as a clear rejection of the RBC paradigm (e.g., Galí, 1999; Galí and Rabanal, 2004). Here, we have shown that dispersed information may potentially help the RBC paradigm accommodate this fact without any need to invoke sticky prices.

Indeed, note that it is the dispersion of information, not the uncertainty about the productivity shock, that causes employment to fall. If agents had been imperfectly informed about the productivity shock but information had been common, then they could fail to increase their employment as much as they would have done with perfect information, but they would not have reduced their employment—for how could they respond to the shock by reducing employment if they were not aware of the shock in the first place? Thus, employment falls in our model precisely because each agent is well informed about the shock but the shock is not common knowledge.

Turning to the effects of noise, in Figure 2 we consider the impulse responses of output and employment in response to a positive innovation in ε_t . As emphasized before, this should be interpreted as a positive error in expectations of aggregate output, rather than as an error in expectations of aggregate fundamentals. When $\alpha = 0$, such forecast errors are irrelevant, simply because individual incentives do not depend on forecasts of aggregate activity. But when $\alpha > 0$, they generate a positive response in output and employment, thus becoming partly self-fulfilling. Furthermore, the stronger the complementarity, the more pronounced the impact of these errors on aggregate employment and output.

The figure considers a positive noise shock, which means a positive shift in expectations about economic activity. The impact of a negative shift in expectations is symmetric. Note that when these shocks occur, output, employment and consumption move in the same direction, without any movement in TFP. The resulting booms and recessions could thus be (mis)interpreted as a certain type of demand shocks. We will return to this point in a moment. Finally, note that the impact of these noise shocks on output and employment can be quite persistent, even though the noise itself is not. This is simply because the associated forecast errors are themselves persistent.

5.2 Variance decomposition and forecast errors

Comparing the responses of employment with those of output to the two shocks, we see that the former is smaller than the latter in the case of productivity shocks but quite larger in the case of

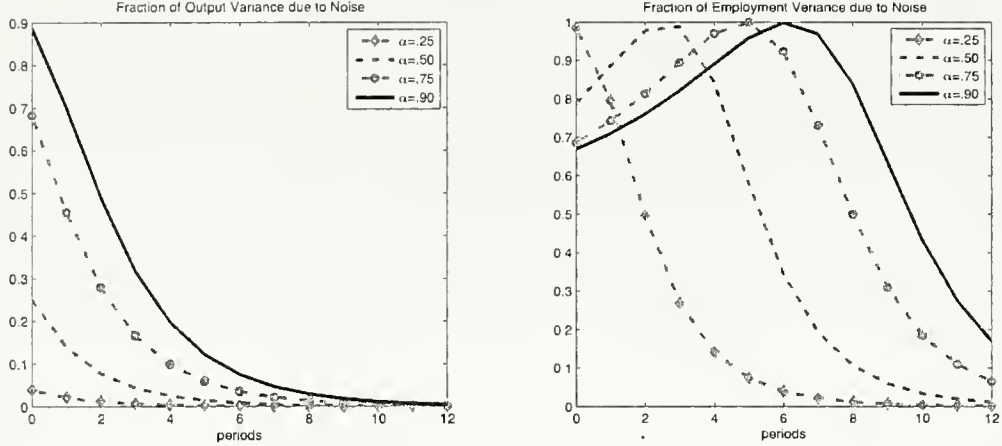


Figure 3: Variance decomposition.

noise. This is simply because productivity shocks have a double effect on output, both directly and indirectly through employment, while the noise impacts output only through employment. But then the response of employment to noise is bound to be stronger than that of output as long as there are diminishing returns to labor ($\theta < 1$), and the more show the lower θ . It follows that noise contributes to a higher relative volatility for employment, while productivity shocks contribute in the opposite direction. In the standard RBC framework, employment may exhibit a higher volatility than output to the extent that there are powerful intertemporal substitution effects (which here we have ruled out since we have also ruled out capital). However, the RBC framework is known to lack in this dimension. Our results here indicate how noise could help improve the performance of the RBC framework in this dimension.

Comparing Figures 1 and 2, it is evident that low-frequency movements in employment and output are dominated by the productivity shocks, while noise contributes relatively more to high-frequency movements. To further illustrate this property, in Figure 3 we plot the variance decomposition of output and employment at different time horizons. For sufficiently strong strategic complementarity, productivity shocks explain only a small fraction of the high-frequency variation in output—short-run fluctuations are driven mostly by noise. As for employment, the contribution of noise is quite dramatic.

Finally, Figure 4 plots the dynamics of the average forecast of aggregate output and the true level of aggregate output in response to a productivity or noise shock. The average forecast error is the distance between the two aforementioned variables. A salient feature of this figure is that forecast errors are *smallest* when the degree of strategic complementarity is highest.

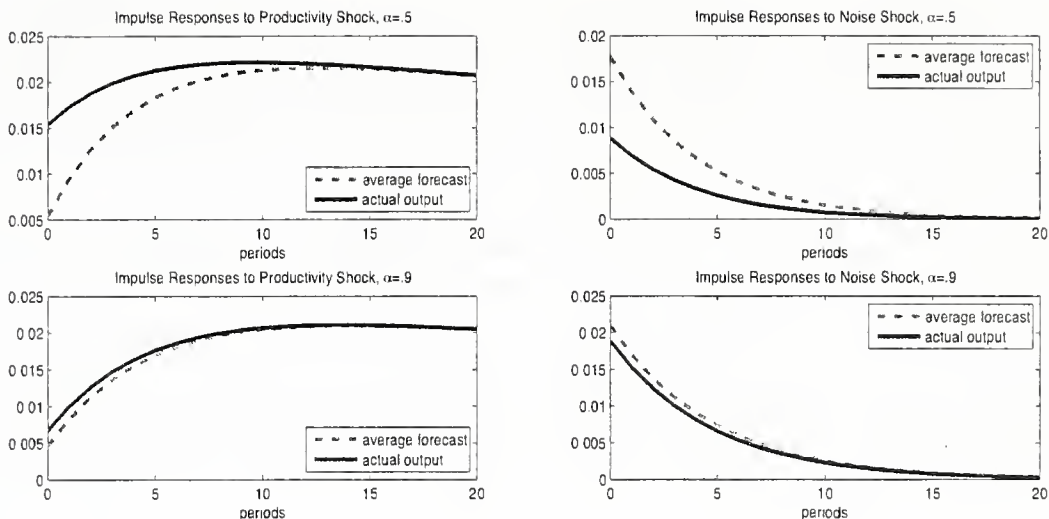


Figure 4: Forecast errors in response to productivity and noise shocks.

This is crucial. We earlier showed that a higher α leads to both more inertia in the response of output and employment to productivity shock, and to a bigger impact of noise. In this sense, the deviation from the common-knowledge benchmark is highest when α is highest. However, one should not expect that these large deviations will show up in large forecast errors. To the contrary, a higher α implies that actual economic activity is more driven by forecasts of economic activity, so that at the end a higher α guarantees that the forecast errors are smaller. It follows that, as we vary α , the magnitude of the deviations of actual outcomes from their common-knowledge counterparts is *inversely* related to the magnitude of the associated forecast errors. Indeed, both the inertia and the impact of noise become nearly self-fulfilling as α gets closer to 1.

Combined, these results illustrate the distinct and powerful mark that the dispersed information can have on macroeconomic outcomes once combined with strong strategic complementarity. Not only can the macroeconomic effects we have documented be quite significant, but they are also consistent with small errors in the agents' forecasts of either the underlying economic fundamentals or the level of economic activity.

5.3 Demand shocks, new-Keynesian models, and structural VARs

Many economists have found the idea that short-run fluctuations are driven primarily by technology shocks implausible either on a priori grounds or on the basis of certain structural VARs. Blanchard and Quah (1989) were the first to attempt to provide some evidence that short-run fluctuations are

driven by “demand” rather than “supply” shocks, albeit with the caveat that one cannot know what the shocks they identify really capture. Subsequent contributions by Galí (1999), Basu, Fernald and Kimball (2006), Galí and Rabanal (2004), and others have tried to improve in that dimension. One way or another, though, this basic view that business cycles are not driven by technology shocks appears to underly the entire New-Keynesian literature.

Our findings here are consistent with this view. In our environment, technology shocks may explain only a small fraction of the high-frequency volatility in macroeconomic outcomes. However, the residual fluctuations have nothing to do with monetary shocks. Rather, they are the product of the noise in the agents’ information. Importantly, to the extent that information is dispersed and trade linkages are important, this noise might be quite small and nevertheless explain a big fraction of the high-frequency volatility in macroeconomic outcomes.

Furthermore, the noise-driven fluctuations we have documented here, albeit being purely neo-classical in their nature, they could well be interpreted as some kind of “demand” or “monetary” shocks in the following sense. This is because they share many of the features often associated with such shocks: they contribute to positive co-movement in employment, output and consumption; they are orthogonal to the underlying productivity shocks; they are closely related to shifts in expectations of aggregate demand; and they explain a large portion of the high-frequency variation in employment and output while vanishing at low frequencies.¹⁸

To better appreciate this, suppose that we generate data from our model using a random-walk specification for the productivity shock and let an applied macroeconomist—preferably of the new-keynesian type—to run a structural VAR as in Blanchard and Quah (1989) or Galí (1999). One would then correctly identify the underlying innovations to productivity by the shock that is allowed to have a long-run effect on output or labor productivity, and the underlying noise shocks by the residual.¹⁹ In the language of Blanchard and Quah, the productivity shocks would be interpreted as “supply shocks” and the noise shocks as “demand shocks”. However, the latter would have no relation to sticky prices and the like; to the contrary, both type of shocks emerge from a purely supply-side mechanism. In the language of Galí (1999), on the other hand, the productivity shocks would be interpreted as “technology shocks”. Furthermore, as already noted, the short-run response of employment to these identified shocks would be negative for high enough α ; but this would not favor a sticky-price interpretation.

¹⁸A priori, our predictions regarding real quantities seems consistent with any predictions about nominal prices. Further exploring under what conditions our noise-driven fluctuations could be associated also with procyclical nominal prices requires a monetary extension of the model.

¹⁹Incidentally, note that the econometric issues studied in Blanchard, L’Huillier, and Lorenzoni (2009) do not apply to the type of noise fluctuations obtained in our model.

As mentioned in the introduction, Beaudry and Portier (2004, 2006), Christiano et al. (2008), Gilchrist and Leahy (2002), Jaimovich and Rebelo (2008), and Lorenzoni (2008) have explored the idea that noisy news about future productivity contribute to short-run fluctuations. Furthermore, Lorenzoni (2008) interprets the resulting fluctuations as “demand shocks” and discusses how they help match related facts. However, all these papers focus on fluctuations that originate from uncertainty about a certain type of fundamentals (namely future productivity), not on the distinct type of uncertainty that emerges when information is heterogeneous and that we highlight in this paper.²⁰ Second, as often the case with new-keynesian models, Lorenzoni’s “demand shocks” confound real shocks with monetary shocks. By this we mean the following. Since there is no capital in his model (as in ours), expectations of future productivity would have been irrelevant for current macroeconomic outcomes had nominal prices been flexible; the only reason then that news about future productivity cause demand-like fluctuations is that they cause an expansion in monetary policy away from the one that would replicate flexible-price allocations. In contrast, our “demand shocks” obtain in an RBC setting and are completely unrelated to monetary policy.

Finally, note that a positive productivity shock in our model induces a small impact on output at high frequencies, followed by a large persistence response at lower frequencies.²¹ Again these properties are consistent with the estimated dynamics of “technology” shocks.

More generally, it is interesting to note that in many empirical new-keynesian models sticky prices dampen the response of output to productivity shocks relative to the RBC framework and help get a negative response for employment. According to some researchers, these properties seem to be more consistent with the data than their RBC counterparts. However, what is a success for these models is only a failure for monetary policy: the only reason that the response of the economy to productivity shocks in the baseline new-keynesian model differs from that in the baseline RBC model is that monetary policy fails to replicate flexible-price allocations, which is typically the optimal thing to do. Here, instead, we obtain the same empirical properties without introducing sticky properties and without presuming any suboptimality for policy.

Perhaps as interestingly, our approach may have intriguing implications for the identification of monetary shocks. One of the standard identification strategies is based on the idea that monetary policy often reacts to measurement error in the level of aggregate economic activity (Bernanke and Mihov, 1995; Christiano, Eichenbaum and Evans, 1999). In particular, the idea is that measurement

²⁰In his baseline model, Lorenzoni considers a representative-agent model with symmetric information. In an extension, he allows for dispersed information, but only to facilitate a more plausible calibration of the model.

²¹In our numerical exercises, the impact of the productivity shock vanishes asymptotically, only because we have assumed that \bar{a}_t is (slowly) mean-reverting. If instead we assume that \bar{a}_t is a random walk, then the long-run impact of a productivity shocks becomes positive, while the rest of the results remain unaffected.

error justifies the existence of random shocks to monetary policy, which are orthogonal to the true underlying state of the economy. If one then traces the impact of these particular shocks on subsequent aggregate outcomes, one can escape the endogeneity problem and identify the impact of monetary shocks. However, these measurement errors, or more generally any forecast errors that the central bank makes about current and future economic activity, are likely to be correlated with the corresponding forecast errors of the private sector. But then the so-identified monetary shocks may actually be proxying for the real effects of the forecast errors of the private sector, which unfortunately are not observed by the econometrician.

5.4 Labor wedges and Solow residuals

Many authors have argued that a good theory of the business cycle must explain the observed variation in the labor wedge and the Solow residual (e.g., Hall, 1997; Rotemberg and Woodford, 1999; Chari, Kehoe, and McGrattan, 2007; Shimer, 2009). We now consider the implications of our model for these two key characteristics of the business cycle.

Following the literature, we define the labor wedge $\tau_{n,t}$ implicitly by

$$\frac{N_t^{\varepsilon-1}}{C_t^{-\gamma}} = (1 - \tau_{n,t}) \theta \frac{Q_t}{N_t}.$$

The left panel of Figure 5 plots the impulse response of the labor wedge to a positive productivity and a positive noise shock. The labor wedge follows very different dynamics in response to the two types of shocks. In particular, a positive productivity shock induces a positive response in the labor wedge, implying positive comovement of the labor wedge with output. On the other hand, a positive noise shock produces a negative response in the observed labor wedge, implying a negative comovement with output.

Multiple authors have documented that variation in the labor wedge plays a large role in accounting for business-cycle fluctuations during the post-war period. Importantly, the labor wedge is highly countercyclical, exhibiting sharp increases during recessions. Shimer (2009) surveys the facts and the multiple explanations that have been proposed for the observed countercyclicality of the labor. These include taxes, shocks to the disutility of labor, mark-up shocks, fluctuations in wage-setting power, and Shimer’s preferred explanation, search frictions in the labor market. Here, we have found that noise offers another possible explanation for the same fact.

We finally consider the potential implications of our results for observed Solow residuals. Towards this goal, we now introduce a variable input in the production function; the optimal use of this input responds to shocks, but is unobserved by the econometrician and is thus absorbed in the Solow residual. As in King and Rebelo (2000), our preferred interpretation of this input is capital

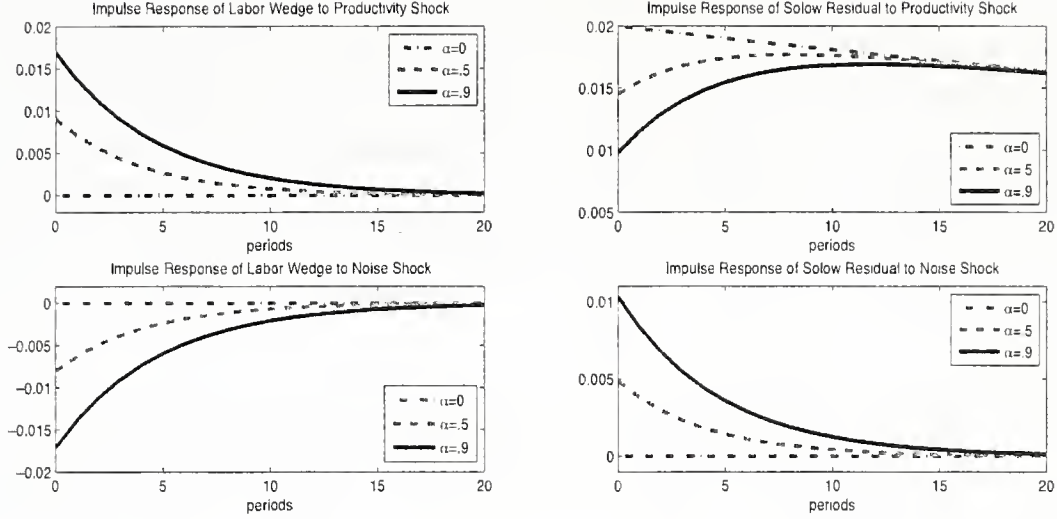


Figure 5: Labor wedges and Solow residuals.

utilization. The only caveat is that in our model capital exogenously fixed. However, we could introduce capital following the same approach as Angeletos and La'O (2009b), without affecting the qualitative points we seek to make here.

We denote the unobserved input by χ_{it} ; we let the gross product of a firm be $\tilde{q}_{it} = \tilde{A}_{it}\chi_{it}^{1-\tilde{\theta}}n_{it}^{\tilde{\theta}}$; and we specify the cost of this input in terms of final product as $\delta\chi_{it}^{1+\xi}$, where $\xi, \delta > 0$. The net product of a firm is then $q_{it} = \tilde{q}_{it} - \delta\chi_{it}^{1+\xi}$. Solving out for the optimal level of this input, The optimal level of this input is given by equating its marginal product with its marginal cost: $(1 - \tilde{\theta})\frac{\tilde{q}_{it}}{\chi_{it}} = \delta(1 + \xi)\chi_{it}^{\xi}$. We thus obtain the following reduced-form production function:

$$q_{it} = A_{it}n_{it}^{\theta} \quad (21)$$

where $\theta \equiv \left(\frac{1+\xi}{\tilde{\theta}+\xi}\right)\tilde{\theta}$ and $A_{it} \equiv \left(\frac{1+\xi}{\tilde{\theta}+\xi}\right)\tilde{A}_{it}^{\frac{1+\xi}{\tilde{\theta}+\xi}}$. Our analysis then remains intact, provided we reinterpret the production function in the above way. Accordingly, we set $\tilde{\theta} = .6$ and $\xi = .1$ (a preferred value in King and Rebelo, 2000), which implies $\theta = .88$. We also re-calibrate the underlying aggregate productivity shocks so that the observed Solow residual ($SR_t \equiv \log Q_t - \theta \log N_t$) implied by the common-knowledge version of the model continues to have a standard deviation of 0.02 and a persistence of 0.99.

The right panel of Figure 5 plots the dynamic response of the Solow residual to a productivity or a noise shock. Both shocks raise the measured Solow residual, but only the innovation in productivity has a persistent effect. Moreover, these responses of the Solow residual mirror those of output. It follows that the Solow residual and output move tightly together, much alike in a standard RBC

model, although employment has the more distinct behavior we mentioned earlier.

Finally, it is worth noting that additional variation in measured Solow residuals could obtain from variation in the dispersion of information, simply because the dispersion of information affects the cross-sectional allocations to resources. Note in particular that the observed heterogeneity in forecast surveys is highly countercyclical, suggesting that the dispersion of information may also be countercyclical. Exploring how such variation in the dispersion of information affects the business cycle is left for future work.

5.5 Discussion

While the characterization of equilibrium in Section 3 allowed for arbitrary information structures, the more concrete positive results that we documented thereafter presumed a specific, Gaussian information structure (Assumption 1). However, we do not expect any of the predictions we have emphasized to be unduly sensitive to the details of the information structure.

We build this expectation on the following observations. Proposition 1 permits us to map our economy to a class of games with linear best responses, like those studied in Morris and Shin (2002) and Angeletos and Pavan (2007, 2009). In this class of games, one can show under arbitrary information structures that a stronger strategic complementarity makes equilibrium outcomes less sensitive to first-order beliefs (the forecasts of the fundamentals) and more sensitive to higher-order beliefs (the forecasts of the forecasts of others). One can then proceed to show quite generally that higher-order beliefs are more sensitive to the initial common prior, to public signals, and to signals with strongly correlated errors, than lower-order beliefs, simply because these pieces of information are relatively better predictors of the forecasts of others. It follows that higher-order beliefs are less sensitive to innovations in the fundamentals and more sensitive to common sources of noise than lower-order beliefs. Combined, these observations explain why stronger complementarity dampens the response of the economy to innovations in fundamentals while amplifying the impact of noise—which are the key properties that drive the results we documented in Sections 4 and 5. We conclude that these results are not unduly sensitive to the details of the underlying information structure; rather, they obtain from robust properties of higher-order beliefs and the very nature of the general-equilibrium interactions in our economy.

Our analysis has implications, not only for aggregate fluctuations, but also for the cross-sectional dispersion of prices and quantities. As evident from condition (24), a higher α necessarily reduces the sensitivity of local output to local fundamentals, while increasing the sensitivity to expectations of aggregate output. When information is commonly shared, all agents share the same expectation of aggregate output, and hence heterogeneity in output (and thereby in prices) can originate *only*

from heterogeneity in fundamentals (productivities, tastes, etc). It then follows that a higher α necessarily reduces cross-sectional dispersion in output and prices, simply because it dampens the only source of heterogeneity. However, once information is dispersed, there is an additional source of heterogeneity: different firms have different expectations of aggregate economic activity. It then follows that a higher α dampens the former source of heterogeneity while amplifying the latter. We conclude that, once information is dispersed, the impact of complementarity on cross-sectional dispersion is ambiguous—which also implies that evidence on the cross-sectional dispersion of prices and quantities may provide little guidance for a quantitative assessment of our results.²²

Similarly, evidence on the size of monopolistic mark-ups, or the elasticity of demands faced by individual firms, do not necessarily discipline the magnitude of our results. This is for two reasons. First, in our model, the mark-up and the elasticity of individual demands identify only η , whereas it is ρ that matters for complementarity. And second, as evident from the definition of α , a high complementarity in our model is consistent with *any* value of ρ , provided that there is a sufficiently small wealth effect on labor supply in the short run, a sufficiently high Frisch elasticity (as in Hansen, 1985), and nearly linear returns to labor in the short run (as in King and Rebelo, 2000).

Finally, it is worth noting that, unlike what suggested by one of our discussants, the findings of Hellwig and Venkateswaran (2009) bear little relevance to our own analysis. That paper considers a New-Keynesian model in which firms cannot tell apart aggregate monetary shocks from idiosyncratic productivity or demand shocks. For a particular calibration, this confusion induces firms to adjust their prices a lot in response to monetary shocks even when these shocks are unobserved, thus leaving little room for real effects. These findings are interesting on their own right—and indeed complement our motivation for focusing on *real* shocks. However, it was unfortunate for our discussant to extrapolate from that paper to the likely quantitative importance of our results. First, the core mechanism of that paper does not apply to our context: if firms were to confuse aggregate shocks for local ones in our model, this confusion would only reinforce our results.²³ And second, the quantitative findings of that paper are based on a number of heroic assumptions, which might serve

²²One of our discussants made the opposite argument. But his argument was based on the premise that a higher α necessarily reduces cross-sectional dispersion. This happens to be true under the specific signal structure we introduced in Assumption 1 but, as just explained, is not true in general. We refer the reader to Angeletos and La'o (2007) for a stark example where higher complementarity increases *both* the non-fundamental component of the business cycle and the cross-sectional dispersion of quantities and prices.

²³To see this, recall from Proposition 3 and Corollary 3 that the response of equilibrium output to an idiosyncratic shock in fundamentals is given by $\varphi_f = 1 - \alpha$, while its response to an aggregate shock is given by $\varphi_v = 1 - \alpha \frac{\kappa_f}{(1-\alpha)\kappa_x + \kappa_y + \kappa_f}$. As long as $\alpha > 0$, φ_f is smaller than φ_v , which means that mistaking an aggregate shock for an idiosyncratic shock only helps dampen the response of the economy to the aggregate shock.

certain purposes but are completely out of place in our own context.²⁴

With these observations we are not trying to escape the need for a serious quantitative exercise, nor are we ready to speculate on the outcome of such an exercise. We are only trying to provide some guidance for any future quantitative exploration of our results. The key effects we have documented in this paper hinge only on (i) the sensitivity of individual output to forecasts of aggregate output and (ii) the sensitivity of these forecasts to the underlying shocks. We are thus skeptical that micro evidence on prices or quantities can alone provide enough guidance on the quantitative importance of our results. We instead propose that a quantitative assessment of our results should rely more heavily on survey evidence about the agents' forecasts of economic activity. Indeed, these forecasts concisely summarize all the informational effects in our model, and their joint stochastic behavior with actual outcomes speaks to the heart of our results.

In this regard, we find the approach taken in Coibion and Gorodnichenko (2008) particularly promising. This paper uses survey evidence to study how the agents' forecasts of certain macroeconomic outcomes respond to certain structural shocks (with the latter being identified by specific structural VARs). In effect, the exercises conducted in that paper are empirical analogues of the theoretical exercise we conducted in Figure 4 for the case of productivity shocks. However, that paper focuses on how the stochastic properties of the forecasts alone could help tell apart specific formalizations of the informational fictions and ignores the fact that the forecasts and actual outcomes are jointly determined. Here, instead, we propose that the emphasis should be shifted from the details of the underlying informational frictions to the joint stochastic properties of these forecasts and the actual macroeconomic outcomes.

²⁴In particular, Hellwig and Venkateswaran (2009) assume that workers are perfectly informed about the monetary shocks, so that nominal wages adjust one-to-one with them. When firms face constant real marginal costs and iso-elastic demands, this assumption can *alone* guarantee that prices will move one-to-one with monetary shocks even if firms cannot tell whether their nominal wages have moved because of nominal or idiosyncratic reasons. Clearly, the empirical relevance of this assumption is questionable even within the context of that paper. As for our own context, we see no good reason for assuming a priori that workers are perfectly informed about the aggregate *real* shocks hitting the economy. Furthermore, Hellwig and Venkateswaran (2009) assume that firms are free to adjust their action at no cost and at a daily or weekly frequency. When that action is interpreted as a nominal price (as in that paper), this assumption serves a useful pedagogical purpose: it helps isolate information frictions from sticky prices. But once that action is interpreted as a real employment or investment choice (as in our model), this assumption stops making sense even from pedagogical perspective: the "stickiness" of real employment and investment decisions is merely a matter of technology in the model.

6 Efficiency

The positive properties we have documented are intriguing. However, their normative content is unclear. Is the potentially high contribution of noise to business-cycle fluctuations, or the potentially high inertia in the response of the economy to innovations in productivity, a symptom of inefficiency?

More generally, it is obvious that a planner could improve welfare if he could centralize all the information that is dispersed in society and then dictate allocations on the basis of all this information. But this would endow the planner with a power that seems far remote from the powers that policy makers have in reality. Furthermore, the resulting superiority of centralized allocations over their decentralized equilibrium counterparts would not be particularly insightful, since it would be driven mostly by the assumption that the planner has the superior power to overcome the information frictions imposed on the market. Thus, following Angeletos and Pavan (2007a, 2009) and Angeletos and La'O (2008), we contend that a more interesting question—on both practical and conceptual grounds—is to understand whether a planner could improve upon the equilibrium while being subject to the same informational frictions as the equilibrium.

This motivates us to consider a constrained efficiency concept that permits the planner to choose any resource-feasible allocation that respects the geographical segmentation of information in the economy—by which we simply mean that the planner cannot make the production and employment choices of firms and workers in one island contingent on the private information of another island. A formal definition of this efficiency concept and a detailed analysis of efficient allocations can be found, for a variant model, in Angeletos and La'O (2008). Here we focus on the essence.

Because of the concavity of preferences and technologies, efficiency dictates symmetry in consumption across households, as well as symmetry across firms and workers within any given island. Using these facts, we can represent the planning problem we are interested in as follows.

Planner's problem. *Choose a pair of local production and employment strategies, $q : \mathcal{S}_\omega \times \mathcal{S}_\Omega \rightarrow \mathbb{R}_+$ and $n : \mathcal{S}_\omega \times \mathcal{S}_\Omega \rightarrow \mathbb{R}_+$, and an aggregate output function, $Q : \mathcal{S}_\Omega^2 \rightarrow \mathbb{R}_+$, so as to maximize*

$$\int_{\mathcal{S}_\Omega} \left[U(Q(\Omega_t, \Omega_{t-1})) - \int_{\mathcal{S}_\omega} \frac{1}{1+\epsilon} S(\omega) n(\omega, \Omega_{t-1})^{1+\epsilon} d\Omega_t(\omega) \right] d\mathcal{P}(\Omega_t | \Omega_{t-1}) \quad (22)$$

subject to

$$q(\omega, \Omega_{t-1}) = A(\omega) n(\omega, \Omega_{t-1})^\theta \quad \forall (\omega, \Omega_{t-1}) \quad (23)$$

$$Q(\Omega_t, \Omega_{t-1}) = \left[\int q(\omega, \Omega_{t-1})^{\frac{\rho-1}{\rho}} d\Omega_t(\omega) \right]^{\frac{\rho}{\rho-1}} \quad \forall (\Omega_t, \Omega_{t-1}) \quad (24)$$

where $\mathcal{P}(\Omega_t | \Omega_{t-1})$ denotes the probability distribution of Ω_t conditional on Ω_{t-1} .

This problem has a simple interpretation. $U(Q(\Omega_t, \Omega_{t-1}))$ is the utility of consumption for the representative household; $\frac{1}{\epsilon} S(\omega) n(\omega, \Omega_{t-1})^\epsilon$ is the marginal disutility of labor for the typical worker in a given island; and the corresponding integral is the overall disutility of labor for the representative household. Furthermore, note that, once the planner picks the production strategy q , the employment strategy n is pinned down by (23) and the aggregate output function Q is pinned down by (23). The reduced-form objective in (22) is thus a functional that gives the level of welfare implied by any arbitrary production strategy that the planner dictates to the economy.

Because this problem is strictly concave, it has a unique solution and this solution is pinned down by the following first-order condition:²⁵

$$S_{it} n_{it}^\epsilon = \mathbb{E}_{it} \left[U'(Q_t) \left(\frac{q_{it}}{Q_t} \right)^{-\frac{1}{\rho}} \right] \left(\theta A_{it} n_{it}^{\theta-1} \right). \quad (25)$$

This condition simply states that the planner dictates the agents to equate the social cost of employment in their island with the *local* expectation of the social value of the marginal product of that employment. Essentially the same condition characterizes (first-best) efficiency in the standard, symmetric-information paradigm. The only difference is that there expectations are conditional on the commonly-available information set, while here they are conditional on the locally-available information sets.

As with equilibrium, we can use $q_{it} = A_{it} n_{it}^\theta$ to eliminate n_{it} in the above condition, thereby reaching the following result.

Proposition 5. *Let*

$$f^*(\omega) \equiv \log \left\{ \theta^{\frac{1}{\frac{\epsilon}{\theta} + \gamma - 1}} \left(\frac{A(\omega)}{S(\omega)} \right)^{\frac{\frac{\epsilon}{\theta}}{\frac{\epsilon}{\theta} + \gamma - 1}} \left(\frac{A(\omega)}{S(\omega)} \right)^{\frac{\frac{\epsilon}{\theta}}{\frac{\epsilon}{\theta} + \gamma - 1}} \right\}$$

be a composite of the local productivity and taste shocks. The efficient strategy $q : \mathcal{S}_\omega \times \mathcal{S}_\Omega \rightarrow \mathbb{R}_+$ is the fixed point to the following:

$$\log q(\omega_t, \Omega_{t-1}) = (1 - \alpha) f^*(\omega_t) + \alpha \log \left\{ \mathbb{E} \left[Q(\Omega_t, \Omega_{t-1})^{\frac{1}{\rho} - \gamma} \middle| \omega_t, \Omega_{t-1} \right]^{\frac{1}{\frac{1}{\rho} - \gamma}} \right\} \quad \forall (\omega_t, \Omega_{t-1}), \quad (26)$$

$$Q(\Omega_t, \Omega_{t-1}) = \left[\int q(\omega, \Omega_{t-1})^{\frac{\rho-1}{\rho}} d\Omega_t(\omega) \right]^{\frac{\rho}{\rho-1}} \quad \forall (\Omega_t, \Omega_{t-1}). \quad (27)$$

A number of remarks are worth making. First, note that the composite shock f_t^* plays a similar role for the efficient allocation as the composite shock f_t played for the equilibrium: it identifies the

²⁵Because of the continuum, the efficient allocation is determined only for *almost* every ω . For expositional simplicity, we bypass the *almost* qualification throughout the paper.

fundamentals that are relevant from the planner's point of view. This is evident, not only from the above result, but also directly from the planner's problem: using $q_t = A_t n_t^\theta$ to eliminate n_t in the expression for welfare given in the planner's problem, we can express welfare as a simple function of the production strategy and the composite shock f_t^* alone.

Second, note that Proposition 5 permits a game-theoretic interpretation of the efficient allocation, much alike what Proposition 1 did for equilibrium: the efficient allocation of the economy coincides with the Bayes-Nash equilibrium of a game in which the different players are the different islands of the economy and their best responses are given by (26).

Third, note that, apart from the different composite shock, the structure of the fixed point that characterizes the efficient and the equilibrium allocation is the same: once we replace $f^*(\omega_t)$ with $f(\omega_t)$, condition (26) coincides with its equilibrium counterpart, condition (7). And because $f^*(\omega_t) = f(\omega_t)$ for every ω_t if and only if there is no monopoly power, the following is immediate.

Corollary 3. *In the absence of monopoly distortions, the equilibrium is efficient, no matter the information structure.*

This result establishes that neither the presence of noise nor the dispersion of information are per se sources of inefficiency. This result might sound bizarre in light of our earlier results that the economy can feature extreme amplification effects, with a tiny amount of noise contributing to large aggregate fluctuations. However, it should be ex post obvious. What causes these large positive effects is the combination of dispersed information and strong complementarity. But neither one introduces a wedge between the equilibrium and the planner. Indeed, the geographical segmentation of information is similar to a technological constraint that impacts equilibrium and efficient allocations in a completely symmetric way. As for the complementarity, it's origin is preferences and technologies, not any type of market inefficiency, guaranteeing that private motives in coordinating economic activity are perfectly aligned with social motives. It follows that, when stronger complementarity amplifies the impact of noise, it does so without causing any inefficiency.²⁶

We can generalize this result for situations where firms have monopoly power, to the extent that there are no aggregate shocks to monopoly power, as follows.

Corollary 4. *Suppose that information is Gaussian (Assumption 1 holds) and there are no aggregate mark-up shocks ($\bar{f}_t^* - \bar{f}_t$ is fixed). Then, the the business cycle is efficient in the sense the gap $\log Q_t - \log Q_t^*$ between the equilibrium and the efficient level of output is invariant.*

If we allow for mark-up shocks, then clearly the equilibrium business cycle ceases to be efficient. But this is true irrespectively of whether information is dispersed or commonly shared. We conclude

²⁶As mentioned earlier, this is the opposite of what happens in Morris and Shin (2002).

that the dispersion of information per se is not a source of inefficiency, whether one considers a competitive RBC or a monopolistic New-Keynesian model. We further discuss the implications of this result for optimal policy and the social value of information in Angeletos and La'O (2008).

We conclude this section with an important qualification. While our efficiency results allowed for an arbitrary information structure, they restricted the information structure to be exogenous to the underlying allocations. This ignores the possibility that information gets endogenously aggregated through prices, macro indicators, and other channels of social learning—which is clearly an important omission. We address this issue, too, in Angeletos and La'O (2008), by allowing information to get partly aggregated through certain price and quantity indicators. We first show that a planner who internalizes the endogeneity of the information contained in these indicators will choose a different allocation than the equilibrium. This typically means that the planner likes to increase the sensitivity of allocations to private information, so as to increase the precision of the information that gets revealed by the available macroeconomic indicators. We then explore policies that could help in this direction.

7 Concluding remarks

The pertinent macroeconomics literature has used informational frictions to motivate why economic agents may happen, or choose, to be partly unaware about the shocks hitting the economy. Sometimes the informational friction is exogenous, sometimes it is endogenized. Invariably, though, the main modeling role of informational frictions seems to remain a simple and basic one: to limit the knowledge that agents have about the underlying shocks to economic fundamentals.

Our approach, instead, seeks to highlight that the *heterogeneity* of information may have a very distinct mark on macroeconomic outcomes than the uncertainty about fundamentals. We highlighted this in this paper by showing how the heterogeneity of information can induce significant inertia in the response of the economy to productivity shocks, and can also generate significant noise-driven fluctuations, even when the agents are well informed about the underlying fundamentals. In Angeletos and La'O (2009b), we further show that the heterogeneity of information can open the door to a novel type of sentiment shocks—namely shocks that are independent of either the underlying fundamentals or the agents' expectations of the fundamentals and nevertheless cause variation in the agents' forecasts of economic activity and thereby in actual economic activity, despite the uniqueness of equilibrium. This in turn permits a broader interpretation of what noise stood for in the present paper: noise could be interpreted more generally as any variation in the forecasts of economic activity that is orthogonal by fundamentals.

In this paper, we focused on the dispersion of information about the *real* shocks hitting the economy, ruling out sticky prices and dismissing any lack of common knowledge about innovations to monetary policy. This, however, does not mean that we see no interesting interaction between dispersed information and nominal frictions. It only means that we find it a good modeling benchmark to assume common knowledge of the current monetary policy. Where we instead see an intriguing interaction between our approach and monetary policy is the following dimension: when there is dispersed information about the underlying *real* shocks hitting the economy and nominal prices are rigid, the response of monetary policy to any information that becomes available about these shocks may be crucial for how the economy responds to these shocks in the first place. This point was first emphasized at a more abstract level by Angeletos and Pavan (2007b, 2009) and is further explored by Angeletos and La'O (2008) and Lorenzoni (2009) within new-Keynesian variants of the economy we have studied in this paper.

We conclude with a comment on the alternative formalizations of informational frictions. For certain questions, one formalization might be preferable to another.²⁷ For other questions, however, the specifics of any particular formalization may prove unnecessary, or even distracting. The results we have emphasized in this paper appear to hinge only on the heterogeneity of information, not on the specific details of the information structure. To highlight this, we showed that the information structure matters for economic outcomes *only* through its impact on the agents' forecasts of aggregate economic activity. We would thus invite other researchers not to commit to any particular formalization of the information structure (including ours), but rather to take a more flexible approach to the modeling of informational frictions. After all, the data cannot possibly inform us about the details of the information structure. What, instead, the data can do is to inform us about the stochastic properties of the agents' forecasts of economic activity—which, as mentioned, is the only channel through which the dispersion of information matters of economic behavior. Thus, in our view, it is only this evidence that should help discipline the theory.

²⁷For example, if one wishes to understand which particular pieces of information agents are likely to pay more attention to, Sims (2003) offers an elegant methodology.

Appendix

Proof of Proposition 1. The characterization of the equilibrium follows directly from the discussion in the main text. Its existence and uniqueness can be obtained by showing that the equilibrium coincides with the solution to a concave planning problem. For the case that there is no monopoly power ($\eta = \infty$), this follows directly from our analysis in Section 6 and in Proposition 5. A similar result can be obtained for the case with monopoly power.

Proof of Proposition 2. This follows from the discussion in the main text.

Proof of Proposition 3. Suppose that, conditional on ω_t and Ω_{t-1} , $Q(\Omega_t, \Omega_{t-1})$ is log-normal, with variance independent of ω_t ; that this is true under the log-normal structure for the underlying shocks and signals we will prove shortly. Using log-normality of Q in condition (7), we infer that the equilibrium production strategy must satisfy condition (9) with

$$const = \frac{\alpha}{2} \left(\frac{1}{\rho} - \gamma \right) \text{Var} [\log Q(\Omega_t, \Omega_{t-1}) | \omega_t, \Omega_{t-1}]$$

and $\text{Var} [\log Q(\Omega_t, \Omega_{t-1}) | \omega_t, \Omega_{t-1}] = \text{Var} [\log Q(\Omega_t, \Omega_{t-1}) | \Omega_{t-1}]$.

We now guess and verify a log-linear equilibrium under the log-normal specification for the shock and information structure. Suppose the equilibrium production strategy takes a log-linear form: $\log q_t = \varphi_0 + \varphi_{-1} \bar{f}_{t-1} + \varphi_f f_t + \varphi_x x_t + \varphi_y y_t$, for some coefficients $(\varphi_{-1}, \varphi_f, \varphi_x, \varphi_y)$. Aggregate output is then given by

$$\log Q(\Omega_t, \Omega_{t-1}) = \varphi'_0 + \varphi_{-1} \bar{f}_{t-1} + (\varphi_f + \varphi_x) \bar{f}_t + \varphi_y y_t \quad (28)$$

where $\varphi'_0 \equiv \varphi_0 + \frac{1}{2} \left(\frac{\rho-1}{\rho} \right) \left[\frac{\varphi_f^2}{\kappa_f} + \frac{\varphi_x^2}{\kappa_x} + 2 \frac{\varphi_f \varphi_x}{\kappa_x} \right]$. It follows that $Q(\Omega_t, \Omega_{t-1})$ is indeed log-normal, with

$$\mathbb{E} [\log Q(\Omega_t, \Omega_{t-1}) | \omega_t, \Omega_{t-1}] = \varphi'_0 + \varphi_{-1} \bar{f}_{t-1} + (\varphi_f + \varphi_x) \mathbb{E} [\bar{f}_t | \omega_t, \Omega_{t-1}] + \varphi_y y_t \quad (29)$$

$$\text{Var} [\log Q(\Omega_t, \Omega_{t-1}) | \omega_t, \Omega_{t-1}] = (\varphi_f + \varphi_x)^2 \left(\frac{1}{\kappa_f + \kappa_x + \kappa_y} \right) \quad (30)$$

where $\mathbb{E} [\bar{f}_t | \omega_t, \Omega_{t-1}] = \frac{\kappa_f}{\kappa_f + \kappa_x + \kappa_y} \psi f_{t-1} + \frac{\kappa_x}{\kappa_f + \kappa_x + \kappa_y} x_t + \frac{\kappa_y}{\kappa_f + \kappa_x + \kappa_y} y_t$. Substituting these expressions into (9) gives us

$$\begin{aligned} \log q(\omega_t, \Omega_{t-1}) &= const + (1 - \alpha) f(\omega) + \alpha (\varphi'_0 + \varphi_{-1} \bar{f}_{t-1} + \varphi_y y_t) \\ &\quad + \alpha (\varphi_f + \varphi_x) \left(\frac{\kappa_f}{\kappa_f + \kappa_x + \kappa_y} \psi f_{t-1} + \frac{\kappa_x}{\kappa_f + \kappa_x + \kappa_y} x_t + \frac{\kappa_y}{\kappa_f + \kappa_x + \kappa_y} y_t \right) \end{aligned}$$

For this to coincide with $\log q(\omega) = \varphi_0 + \varphi_{-1}\bar{f}_{t-1} + \varphi_f f + \varphi_x x + \varphi_y y$ for every (f, x, y) , it is necessary and sufficient that the coefficients $(\varphi_0, \varphi_{-1}, \varphi_f, \varphi_x, \varphi_y)$ solve the following system:

$$\begin{aligned}\varphi_0 &= \text{const} + \alpha\varphi'_0 \\ \varphi_f &= 1 - \alpha \\ \varphi_x &= \alpha(\varphi_f + \varphi_x) \left(\frac{\kappa_x}{\kappa_f + \kappa_x + \kappa_y} \right) \\ \varphi_{-1} &= \alpha\varphi_{-1} + \alpha(\varphi_f + \varphi_x) \left(\frac{\kappa_f}{\kappa_f + \kappa_x + \kappa_y} \right) \psi \\ \varphi_y &= \alpha\varphi_y + \alpha(\varphi_f + \varphi_x) \left(\frac{\kappa_y}{\kappa_f + \kappa_x + \kappa_y} \right)\end{aligned}$$

The unique solution to this system for $(\varphi_{-1}, \varphi_f, \varphi_x, \varphi_y)$ is the one given in the proposition; φ_0 is then uniquely determined from the first equation of this system along with the definition of *const* and φ'_0 .

Proof of Proposition 4. The result follows by a triple limit. First, take $\alpha \rightarrow 1$; next, take $\kappa_y \rightarrow 0$; and finally, take $\kappa_x \rightarrow \infty$. It is easy to check that this triple limit implies $\kappa \rightarrow \infty$ and $R \rightarrow 1$. That is, the precision of the agents posterior about the fundamentals (the mean squared forecast error) converges to zero, while the fraction of the high-frequency variation in output that is due to noise converges to 100%.

Kalman filtering for dynamic extension. The method we use in solving this equilibrium is similar to that found in Woodford (2003b).

State Vector and Law of Motion. We guess and verify that the relevant aggregate state variables of the economy at time t are \bar{f}_t and $\log Q_t$ and thus define state vector X_t in (18) accordingly.

Claim. *The dynamics of the economy are given by the following law of motion*

$$X_t = MX_{t-1} + m_v v_t + m_\varepsilon \varepsilon_t \quad (31)$$

with

$$M \equiv \begin{bmatrix} \psi & 0 \\ M_{21} & M_{22} \end{bmatrix}, m_v \equiv \begin{bmatrix} 1 \\ m_{v2} \end{bmatrix}, m_\varepsilon \equiv \begin{bmatrix} 0 \\ m_{\varepsilon 2} \end{bmatrix}. \quad (32)$$

The coefficients $(M_{21}, M_{22}, m_{v2}, m_{\varepsilon 2})$ are given by

$$M_{21} = \psi(K_{21} + K_{22}) \quad (33)$$

$$M_{22} = \psi(1 - K_{21} - K_{22}) \quad (34)$$

$$m_{v2} = 1 - \alpha(1 - K_{21} - K_{22}) \quad (35)$$

$$m_{\varepsilon 2} = \alpha K_{22} \quad (36)$$

and

$$K \equiv \begin{bmatrix} K_{11} & K_{21} \\ K_{21} & K_{22} \end{bmatrix}$$

is the matrix of kalman gains, defined by

$$K \equiv \mathbb{E} [(X_t - \mathbb{E}_{i,t-1} [X_t]) (z_{i,t} - \mathbb{E}_{i,t-1} [z_{i,t}])'] \mathbb{E} [(z_{i,t} - \mathbb{E}_{i,t-1} [z_{i,t}]) (z_{i,t} - \mathbb{E}_{i,t-1} [z_{i,t}])']^{-1} \quad (37)$$

We verify this claim in the following and describe the procedure for finding the fixed point.

Observation Equation. In each period t , firms and workers on island i observe vector $z_{i,t}$, as in (19), of private and public signals. In terms of the aggregate state and error terms, island i 's observation equation takes the form

$$z_{i,t} \equiv \begin{bmatrix} e'_1 \\ e'_1 \end{bmatrix} X_t + \begin{bmatrix} 1 \\ 0 \end{bmatrix} \varsigma_{it} + \begin{bmatrix} 0 \\ 1 \end{bmatrix} \varepsilon_t \quad (38)$$

where e_j is defined as a column vector of length two where the j -th entry is 1 and all other entries are 0.

Forecasting and Inference. Island i 's $t - 1$ forecast of z_t^i is given by

$$\mathbb{E}_{i,t-1} [z_{i,t}] = \begin{bmatrix} e'_1 \\ e'_1 \end{bmatrix} \mathbb{E}_{i,t-1} [X_t]$$

where $\mathbb{E}_{i,t-1} [X_t]$ is island's i 's $t - 1$ forecast of X_t . Combining this with the law of motion (31), it follows that $\mathbb{E}_{i,t-1} [X_t] = M \mathbb{E}_{i,t-1} [X_{t-1}]$.

To form minimum mean-squared-error estimates of the current state, firms and workers on each island use the kalman filter to update their forecasts. Updating is done via

$$\mathbb{E}_{i,t} [X_t] = \mathbb{E}_{i,t-1} [X_t] + K (z_{i,t} - \mathbb{E}_{i,t-1} [z_{i,t}]), \quad (39)$$

where K is the 2×2 matrix of Kalman gains, defined in (37). Substitution of island i 's $t - 1$ forecast of z_t^i into (39) gives us

$$\mathbb{E}_{i,t} [X_t] = \left(I - K \begin{bmatrix} e'_1 \\ e'_1 \end{bmatrix} \right) M \mathbb{E}_{i,t-1} [X_{t-1}] + K z_{i,t} \quad (40)$$

Let $\bar{\mathbb{E}}_t [X_t] \equiv \int \mathbb{E}_{i,t} [X_t] di$ be the time t average expectation of the current state. Aggregation over (40) implies

$$\bar{\mathbb{E}}_t [X_t] = \left(I - K \begin{bmatrix} e'_1 \\ e'_1 \end{bmatrix} \right) M \bar{\mathbb{E}}_{t-1} [X_{t-1}] + K \int z_{i,t} di$$

Finally, using the fact that aggregation over signals yields $\int z_{i,t} di = \begin{bmatrix} e'_1 \\ e'_1 \end{bmatrix} X_t + \begin{bmatrix} 0 \\ 1 \end{bmatrix} \varepsilon_t$, it follows that the average expectation evolves according to

$$\begin{aligned} \bar{\mathbb{E}}_t[X_t] &= K \begin{bmatrix} e'_1 \\ e'_1 \end{bmatrix} M X_{t-1} + \left(I - K \begin{bmatrix} e'_1 \\ e'_1 \end{bmatrix} \right) M \bar{\mathbb{E}}_{t-1}[X_{t-1}] \\ &\quad + K \begin{bmatrix} e'_1 \\ e'_1 \end{bmatrix} m_v v_t + K \left(\begin{bmatrix} e'_1 \\ e'_1 \end{bmatrix} m_\varepsilon + \begin{bmatrix} 0 \\ 1 \end{bmatrix} \right) \varepsilon_t \end{aligned} \quad (41)$$

where M, m_v, m_ε are given by (32).

Characterizing Aggregate Output. Local output in each island is determined by the best-response-like condition in (17), which may be rewritten as $\log q_{i,t} = (1 - \alpha) f_t + \alpha e'_{2,t} \mathbb{E}_{i,t}[X_t]$. Aggregating over this condition, we find that aggregate output must satisfy

$$\log Q_t = (1 - \alpha) \bar{f}_t + \alpha e'_{2,t} \bar{\mathbb{E}}_t[X_t] \quad (42)$$

Substituting our expression for $\bar{\mathbb{E}}_t[X_t]$ from (41) into (42), gives us

$$\begin{aligned} \log Q_t &= [(1 - \alpha) \psi + \alpha \psi (K_{21} + K_{22})] \bar{f}_{t-1} + [\alpha M_{21} - \alpha \psi (K_{21} + K_{22})] \bar{\mathbb{E}}_{t-1}[\bar{f}_{t-1}] \\ &\quad + \alpha M_{22} \bar{\mathbb{E}}_{t-1}[\log Q_{t-1}] + [(1 - \alpha) + \alpha (K_{21} + K_{22})] v_t + \alpha K_{22} \varepsilon_t \end{aligned}$$

Moreover, rearranging condition (42), we find that $\bar{\mathbb{E}}_t[\log Q_t] = \frac{1}{\alpha} (\log Q_t - (1 - \alpha) \bar{f}_t)$. Finally, using this condition in the above equation gives us

$$\begin{aligned} \log Q_t &= [(1 - \alpha) \psi + \alpha \psi (K_{21} + K_{22}) - M_{22} (1 - \alpha)] \bar{f}_{t-1} + M_{22} \log Q_{t-1} \\ &\quad + [\alpha M_{21} - \alpha \psi (K_{21} + K_{22})] \bar{\mathbb{E}}_{t-1}[\bar{f}_{t-1}] + [1 - \alpha + \alpha (K_{21} + K_{22})] v_t + \alpha K_{22} \varepsilon_t \end{aligned}$$

For this to coincide with the law of motion conjectured in (31) and (32) for every $(\bar{f}_{t-1}, \log Q_{t-1}, v_t, \varepsilon_t)$, it is necessary and sufficient that the coefficients $(M_{21}, M_{22}, m_{v2}, m_{\varepsilon2})$ solve the following system:

$$\begin{aligned} M_{21} &= (1 - \alpha) \psi + \alpha \psi (K_{21} + K_{22}) - M_{22} (1 - \alpha) \\ m_{v2} &= 1 - \alpha + \alpha (K_{21} + K_{22}) \\ m_{\varepsilon2} &= \alpha K_{22} \\ 0 &= \alpha M_{21} - \alpha \psi (K_{21} + K_{22}) \end{aligned}$$

The unique solution to this system for $(M_{21}, M_{22}, m_{v2}, m_{\varepsilon2})$ is the one given in the proposition. Therefore, given the kalman gains matrix K , we can uniquely identify the coefficients of the law of motion of X_t .

Kalman Filtering. Let us define the variance-covariance matrices of forecast errors as

$$\begin{aligned}\Sigma &\equiv \mathbb{E}[(X_t - \mathbb{E}_{i,t-1}[X_t])(X_t - \mathbb{E}_{i,t-1}[X_t])'] \\ V &\equiv \mathbb{E}[(X_t - \mathbb{E}_{i,t}[X_t])(X_t - \mathbb{E}_{i,t}[X_t])']\end{aligned}$$

These matrices will be the same for all islands i , since their observation errors are assumed to have the same stochastic properties. Using these matrices, we may write K as the product of two components:

$$\mathbb{E}_i[(X_t - \mathbb{E}_{i,t-1}[X_t])(z_{i,t} - \mathbb{E}_{i,t-1}[z_{i,t}])'] = \Sigma \begin{bmatrix} e_1 & e_1 \end{bmatrix} + \sigma_\varepsilon^2 m_\varepsilon \begin{bmatrix} 0 & 1 \end{bmatrix}$$

and

$$\begin{aligned}\mathbb{E}_i[(z_{i,t} - \mathbb{E}_{i,t-1}[z_{i,t}])(z_{i,t} - \mathbb{E}_{i,t-1}[z_{i,t}])'] &= \begin{bmatrix} e'_1 \\ e'_1 \end{bmatrix} \Sigma \begin{bmatrix} e_1 & e_1 \end{bmatrix} + \sigma_v^2 \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \\ &+ \sigma_\varepsilon^2 \left(\begin{bmatrix} e'_1 \\ e'_1 \end{bmatrix} m_\varepsilon \begin{bmatrix} 0 & 1 \end{bmatrix} + \begin{bmatrix} 0 \\ 1 \end{bmatrix} m'_\varepsilon \begin{bmatrix} e_1 & e_1 \end{bmatrix} + \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix} \right)\end{aligned}\quad (43)$$

Therefore, K is given by

$$K = \left(\Sigma \begin{bmatrix} e_1 & e_1 \end{bmatrix} + \sigma_\varepsilon^2 m_\varepsilon \begin{bmatrix} 0 & 1 \end{bmatrix} \right) (\sigma_z^2)^{-1} \quad (44)$$

where $\sigma_z^2 \equiv \mathbb{E}_i[(z_{i,t} - \mathbb{E}_{i,t-1}[z_{i,t}])(z_{i,t} - \mathbb{E}_{i,t-1}[z_{i,t}])']$ is given by (43).

Finally, what remains to determine is the matrix Σ . The law of motion implies that matrices Σ and V satisfy

$$\Sigma = MVM' + \sigma_v^2 m_v m'_v + \sigma_\varepsilon^2 m_\varepsilon m'_\varepsilon,$$

In addition, the forecasting equation (40) imply these matrices must further satisfy

$$V = \Sigma - \left(\Sigma \begin{bmatrix} e_1 & e_1 \end{bmatrix} + \sigma_\varepsilon^2 m_\varepsilon \begin{bmatrix} 0 & 1 \end{bmatrix} \right) (\sigma_z^2)^{-1} \left(\begin{bmatrix} e'_1 \\ e'_1 \end{bmatrix} \Sigma + \sigma_\varepsilon^2 \begin{bmatrix} 0 \\ 1 \end{bmatrix} m'_\varepsilon \right)$$

Combining the above two equations, we obtain the stationary *Ricatti Equation* for Σ :

$$\begin{aligned}\Sigma &= M\Sigma M' - M \left(\Sigma \begin{bmatrix} e_1 & e_1 \end{bmatrix} + \sigma_\varepsilon^2 m_\varepsilon \begin{bmatrix} 0 & 1 \end{bmatrix} \right) (\sigma_z^2)^{-1} \left(\begin{bmatrix} e'_1 \\ e'_1 \end{bmatrix} \Sigma + \sigma_\varepsilon^2 \begin{bmatrix} 0 \\ 1 \end{bmatrix} m'_\varepsilon \right) M' \\ &+ \sigma_v^2 m_v m'_v + \sigma_\varepsilon^2 m_\varepsilon m'_\varepsilon\end{aligned}\quad (45)$$

where M , m_v , m_ε are functions of the kalman gains matrix K , and K is itself a function of Σ and m_ε . The variance-covariance matrix Σ , the kalman gains matrix K , and the law of motion matrices M , m_v , m_ε are thus obtained by solving the large non-linear system of equations described by (33)-(36), (44), and (45). This system is too complicated to allow further analytical results; we thus solve for the fixed point numerically.

Proof of Proposition 5. The planner's problem is strictly convex, guaranteeing that its solution is unique and is pinned down by its first-order conditions. The Lagrangian of this problem can be written as

$$\begin{aligned} \Lambda = & \int_{S_\Omega} \left[U(Q(\Omega_t, \Omega_{t-1})) - \int_{S_\omega} \frac{1}{1+\epsilon} S(\omega) e^{-\frac{1+\epsilon}{\theta} a} q(\omega, \Omega_{t-1})^{\frac{1+\epsilon}{\theta}} d\Omega_t(\omega) \right] d\mathcal{F}(\Omega_t | \Omega_{t-1}) \\ & + \int_{S_\Omega} \lambda(\Omega_t) \left[Q(\Omega_t, \Omega_{t-1})^{\frac{\rho-1}{\rho}} - \int_{S_\omega} q(\omega, \Omega_{t-1})^{\frac{\rho-1}{\rho}} d\Omega_t(\omega) \right] d\mathcal{F}(\Omega_t | \Omega_{t-1}) \end{aligned}$$

The first-order conditions with respect to $Q(\Omega)$ and $q(\omega)$ are given by the following:

$$U'(Q(\Omega_t, \Omega_{t-1})) + \lambda(\Omega_t) \left(\frac{\rho-1}{\rho} \right) Q(\Omega_t, \Omega_{t-1})^{-\frac{1}{\rho}} = 0 \quad (46)$$

$$\int_{S_\Omega} \left[-\frac{1}{\theta} S(\omega) e^{-\frac{1+\epsilon}{\theta} a} q(\omega, \Omega_{t-1})^{\frac{1+\epsilon}{\theta}-1} - \lambda(\Omega_t) \left(\frac{\rho-1}{\rho} \right) q(\omega, \Omega_{t-1})^{-\frac{1}{\rho}} \right] \mathcal{F}(\Omega_t | \omega, \Omega_{t-1}) = 0 \quad (47)$$

where $\mathcal{F}(\Omega_t | \omega, \Omega_{t-1})$ denotes the posterior about Ω_t (or, equivalently, about \bar{f}_t and y_t) given ω_t . Restating condition (46) as $\lambda(\Omega_t) \left(\frac{\rho-1}{\rho} \right) = -U'(Q(\Omega_t, \Omega_{t-1})) Q(\Omega_t, \Omega_{t-1})^{\frac{1}{\rho}}$ and substituting this into condition (47), gives condition (26), which concludes the proof.

References

- [1] Adam, Klaus (2007), "Optimal monetary policy with imperfect common knowledge?," *Journal of Monetary Economics* 54, 276-301.
- [2] Amador, Manuel, and Pierre-Olivier Weill (2007), "Learning from Private and Public Observations of Others' Actions," Stanford University/UCLA mimeo.
- [3] Amador, Manuel, and Pierre-Olivier Weill (2008), "Learning from Prices: Public Communication and Welfare," Stanford University/UCLA mimeo.
- [4] Amato, Jeffery, and Hyun Song Shin (2006), "Imperfect Common Knowledge and the Information Value of Prices," *Economic Theory* 27, 213-241.
- [5] Angeletos, George-Marios, and Jennifer La'O (2009a), "Incomplete Information, Higher-Order Beliefs, and Price Inertia," MIT mimeo.
- [6] Angeletos, George-Marios, and Jennifer La'O (2009b), "Sentiments," MIT mimeo.
- [7] Angeletos, George-Marios, and Jennifer La'O (2008), "Dispersed Information over the Business Cycle: Optimal Fiscal and Monetary Policy with Dispersed Information," MIT mimeo.
- [8] Angeletos, George-Marios, Guido Lorenzoni, and Alessandro Pavan (2008), "Wall Street and Silicon Valley: A Delicate Interaction," MIT/Northwestern mimeo.

- [9] Angeletos, George-Marios, and Alessandro Pavan (2009), "Policy with Dispersed Information," *Journal of the European Economic Association* 7, 11-60.
- [10] Angeletos, George-Marios, and Alessandro Pavan (2007a), "Efficient Use of Information and Social Value of Information," *Econometrica* 75:4, 1103-1142.
- [11] Angeletos, George-Marios, and Alessandro Pavan (2007b), "Socially Optimal Coordination: Characterization and Policy Implications," *Journal of the European Economic Association* 5, 585-593.
- [12] Angeletos, George-Marios, and Alessandro Pavan (2004), "Transparency of Information and Coordination in Economies with Investment Complementarities," *American Economic Review* 94, 91-98.
- [13] Bacchetta, Philippe, and Eric van Wincoop (2005), "Can Information Heterogeneity Explain the Exchange Rate Determination Puzzle?," *American Economic Review* 95.
- [14] Barro, Robert (1976), "Rational Expectations and the Role of Monetary Policy," *Journal of Monetary Economics* 2, 1-32.
- [15] Basu, Susanto, John Fernald, and Miles Kimball (2006), "Are Technology Improvements Contractionary?," *American Economic Review* 96, 1418-1448.
- [16] Beaudry, Paul, and Franck Portier (2004), "Exploring Pigou's Theory of Cycles," *Journal of Monetary Economics* 51, 1183-1216.
- [17] Beaudry, Paul, and Franck Portier (2006), "Stock Prices, News, and Economic Fluctuations," *American Economic Review* 96, 1293-1307.
- [18] Benhabib, Jess, Richard Rogerson, and Randall Wright (1991), "Homework in Macroeconomics: Household Production and Aggregate Fluctuations," *Journal of Political Economy* 99, 1166-1197.
- [19] Bernanke, Ben, and Ilian Mihov (1995), "Measuring Monetary Policy," *The Quarterly Journal of Economics* 110, 869-902.
- [20] Blanchard, Olivier, Jean-Paul L'Huillier, and Guido Lorenzoni (2009), "News, Noise, and Fluctuations: An empirical exploration," MIT mimeo.
- [21] Blanchard, Olivier, and Danny Quah (1989), "The Dynamic Effects of Aggregate Demand and Supply Disturbances," *American Economic Review* 79, 655-673.
- [22] Chari, V. V., Patrick Kehoe, and Ellen McGrattan (2007), "Business Cycle Accounting," *Econometrica* 75, 781-836.
- [23] Christiano, Lawrence, Martin Eichenbaum, and Charles Evans (1999), "Monetary Policy Shocks: What Have We Learned and to What End?," in J. B. Taylor & M. Woodford, eds., *Handbook of Macroeconomics*, Elsevier.
- [24] Christiano, Lawrence, Martin Eichenbaum, and Robert Vigfusson (2003), "What Happens After A Technology Shock?" Board of Governors of the Federal Reserve System, International Finance Discussion Papers No. 768.

- [25] Christiano, Lawrence, Cosmin Ilut, Roberto Motto, and Massimo Rostagno (2008), "Monetary policy and stock market boom-bust cycles," ECB Working Paper No. 955.
- [26] Coibion, Olivier, and Yuriy Gorodnichenko (2008), "What can survey forecasts tell us about informational rigidities?" UC Berkeley mimeo.
- [27] Collard, Fabrice, and Harris Dellas (2005), "Misperceived vs Unanticipated Money: A Synthesis," CNRS-GREMAQ/University of Bern mimeo.
- [28] Galí, Jordi (1999), "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?," *American Economic Review* 89, 249-271.
- [29] Galí, Jordi (2003), "New Perspectives on Monetary Policy, Inflation, and the Business Cycle," in M. Dewatripont, L. Hansen, and S. Turnovsky (eds.), *Advances in Economic Theory*, vol. III, 151-197, Cambridge University Press.
- [30] Galí, Jordi, and Pau Rabanal (2004), "Technology Shocks and Aggregate Fluctuations: How Well Does the Real Business Cycle Model Fit Postwar U.S. Data?," *NBER Macroeconomics* 2004.
- [31] Gilchrist, Simon, and John Leahy (2002), "Monetary Policy and Asset Prices," *Journal of Monetary Economics* 49, 75-97.
- [32] Goodfriend, Marvin, and Robert King (2001), "The Case for Price Stability," NBER Working Paper 8423.
- [33] Goodfriend, Marvin, and Robert King (1997), "The New Neoclassical Synthesis and the Role of Monetary Policy," *NBER Macroeconomics Annual* 1997, Cambridge, MIT Press.
- [34] Hall, Robert E. (1997), "Macroeconomic Fluctuations and the Allocation of Time," *Journal of Labor Economics*, 15, S223-S250.
- [35] Hansen, Gary (1985), "Indivisible Labor and the Business Cycle," *Journal of Monetary Economics* 16, 309-337.
- [36] Hellwig, Christian (2002), "Public Announcements, Adjustment Delays, and the Business Cycle," UCLA mimeo.
- [37] Hellwig, Christian (2005), "Heterogeneous Information and the Welfare Effects of Public Information Disclosures," UCLA mimeo.
- [38] Hellwig, Christian, and Laura Veldkamp (2008), "Knowing What Others Know: Coordination Motives in Information Acquisition," *Review of Economic Studies*, forthcoming.
- [39] Hellwig, Christian, and Venky Venkateswara (2008), "Setting The Right Prices For The Wrong Reasons," UCLA mimeo.
- [40] Jaimovich, Nir, and Sergio Rebelo (2006), "Can News About the Future Drive the Business Cycle?," Stanford/Northwestern mimeo.

- [41] King, Robert (1982), "Monetary Policy and the Information Content of Prices," *Journal of Political Economy* 90, 247-279.
- [42] King, Robert and Sergio Rebelo (2000), "Resusitating Real Business Cycles," in J. B. Taylor & M. Woodford, eds., *Handbook of Macroeconomics*, Elsevier.
- [43] Kydland, Finn, and Edward Prescott (1982), "Time to Build and Aggregate Fluctuations," *Econometrica* 50, 1345-70.
- [44] Lorenzoni, Guido (2008), "A Theory of Demand Shocks," *American Economic Review*, forthcoming.
- [45] Lorenzoni, Guido (2009), "Optimal Monetary Policy with Uncertain Fundamentals and Dispersed Information," *Review of Economic Studies*, forthcoming.
- [46] Lucas, Robert E., Jr. (1972), "Expectations and the Neutrality of Money," *Journal of Economic Theory* 4, 103-124.
- [47] Lucas, Robert E., Jr. (1975), "An Equilibrium Model of the Business Cycle," *Journal of Political Economy* 83, 1113-1144.
- [48] Luo, Yulei (2008), "Consumption Dynamics under Information Processing Constraints," *Review of Economic Dynamics* 11, 366-385.
- [49] Mackowiak, Bartosz, and Mirko Wiederholt (2008), "Optimal Sticky Prices under Rational Inattention," *American Economic Review*, forthcoming.
- [50] Mackowiak, Bartosz, and Mirko Wiederholt (2009), "Business Cycle Dynamics under Rational Inattention," European Central Bank/Nothwestern University mimeo.
- [51] Mankiw, N. Gregory and Ricardo Reis (2002), "Sticky Information Versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve," *Quarterly Journal of Economics* 117, 1295-1328.
- [52] Mankiw, N. Gregory and Ricardo Reis (2006), "Pervasive Stickiness," *American Economic Review* 96, 164-169.
- [53] McGrattan, Ellen (2004), "Comment on Galí and Rabanal," Federal Reserve Bank of Minneapolis, Research Department Staff Report 338.
- [54] Moscarini, Giuseppe (2004), "Limited Information Capacity as a Source of Inertia," *Journal of Economic Dynamics and Control* 28, 2003-2035.
- [55] Morris, Stephen, and Hyun Song Shin (2002), "The Social Value of Public Information," *American Economic Review* 92, 1521-1534.
- [56] Morris, Stephen, and Hyun Song Shin (2006), "The Inertia of Forward-Looking Expectations," *American Economic Review* 96, 1521-1534.
- [57] Nimark, Kristoffer (2008), "Dynamic Pricing and Imperfect Common Knowledge," *Journal of Monetary Economics* 55, 365-382.

- [58] Phelps, Edmund (1970), "Introduction: The New Microeconomics in Employment and Inflation Theory," in A.A. Alchian, C.C. Holt et al., *Microeconomic Foundations of Employment and Inflation Theory*, New York: Norton.
- [59] Prescott, Edward (1986), "Theory Ahead of Business Cycle Measurement," *Federal Reserve Bank Minneapolis Quarterly Review* 10, 9-22.
- [60] Reis, Ricardo (2006), "Inattentive Producers," *Review of Economic Studies* 73, 793-821.
- [61] Reis, Ricardo (2008), "Optimal Monetary Policy Rules in an Estimated Sticky-Information Model," *American Economic Journal: Macroeconomics*, forthcoming.
- [62] Rodina, Giacomo (2008), "Incomplete Information and Informative Pricing," UCSD mimeo.
- [63] Rogerson, Richard (1988), "Indivisible Labor, Lotteries and Equilibrium," *Journal of Monetary Economics* 21, 3-16.
- [64] Rotemberg, Julio, and Michael Woodford (1991), "Markups and the Business Cycle," in O.J. Blanchard and S. Fischer (eds.), *NBER Macroeconomics Annual 1991*, Cambridge, MIT Press.
- [65] Rotemberg, Julio, and Michael Woodford (1997), "An Optimization-Based Econometric Model for the Evaluation of Monetary Policy," in O.J. Blanchard and S. Fischer (eds.), *NBER Macroeconomics Annual 1997*, Cambridge, MIT Press.
- [66] Rotemberg, Julio, and Michael Woodford (1999), "The Cyclical Behavior of Prices and Costs," in J. B. Taylor & M. Woodford (eds.), *Handbook of Macroeconomics*, Elsevier.
- [67] Shimer, Robert (2009), "Convergence in Macroeconomics: The Labor Wedge," *American Economic Journal: Macroeconomics* 1, 280-297.
- [68] Sims, Christopher (2003), "Implications of Rational Inattention," *Journal of Monetary Economics* 50, 665-690.
- [69] Sims, Christopher (2006), "Rational Inattention: Beyond the Linear-Quadratic Case," *American Economic Review* 96, 158-163.
- [70] Townsend, Robert M. (1983), "Forecasting the Forecasts of Others," *Journal of Political Economy* 91, 546-588.
- [71] Van Nieuwerburgh, Stijn, and Laura Veldkamp (2008), "Information Acquisition and Under-Diversification," NYU Stern mimeo.
- [72] Van Nieuwerburgh, Stijn, and Laura Veldkamp (2006), "Learning Asymmetries in Real Business Cycles," *Journal of Monetary Economics* 53, 753-772.
- [73] Veldkamp, Laura (2006), "Media Frenzies in Markets for Financial Information," *American Economic Review* 96, 577-601.

- [74] Veldkamp, Laura, and Justin Woolfers (2006), "Aggregate Shocks or Aggregate Information? Costly Information and Business Cycle Comovement," *Journal of Monetary Economics* 54, 37-55.
- [75] Woodford, Michael (2003a), "Imperfect Common Knowledge and the Effects of Monetary Policy," in P. Aghion, R. Frydman, J. Stiglitz, and M. Woodford, eds., *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*, Princeton University Press.
- [76] Woodford, Michael (2003b), *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton: Princeton University Press.
- [77] Woodford, Michael (2008), "Information-Constrained State Dependent Pricing," Columbia mimeo.